

MULTIFRACTAL AND GAUSSIAN FRACTIONAL SUM-DIFFERENCE MODELS FOR INTERNET TRAFFIC

BY DAVID ANDERSON

BY WILLIAM S. CLEVELAND*

BY BOWEI XI†

A multifractal fractional sum-difference model (MFSD) is a monotone transformation of a Gaussian fractional sum-difference model (GFSD). The GFSD is the sum of two independent components: a moving sum of length two of discrete fractional Gaussian noise (fGn); and white noise. Internet packet traffic interarrival times are very well modeled by an MFSD in which the marginal distribution is Weibull; this is validated by extensive model checking for 715,665,213 measured arrival times on three Internet links. The simplicity of the model provides a mathematical tractability that results in much insight into traffic statistical properties. In the past, the foundation for understanding the properties has been changes in the properties with the time scale; this is a frequency domain foundation. Emerging from mathematical investigations based on the MFSD is a more fundamental foundation based on how the fGn and white noise components, and their relative variances, affect changes in the statistics with changing factors such as the packet arrival rate and time aggregation of the traffic; this is a time domain foundation. A simple logistic model relates the MFSD model parameters to the packet rate. This enables the MFSD model to be used, with just a specification of the rate, to generate packet arrivals for simulation studies.

1. Introduction.

1.1. *Internet Technology.* Internet traffic results from the transfers of information between pairs of computers, or hosts, across the Internet ([Kesidis, 2007](#); [Peterson and Davie, 1999](#); [Stevens, 1994](#)). For simplicity we will refer to the information as a file. The file is broken up into packets with sizes typically up to 1460 bytes = 11680 bits. The packets are sent from the source host over a path consisting of routers connected by transmission links, and the file is reassembled at the destination host. The two hosts establish a connection to carry out the transfer, which means each is listening for the arrival of packets from the other. Headers,

*Communicating author. Supported by ARO MURI Award W911NF-08-1-0238 and NSF Award 09-525.

†Supported by NSF DMS-0904548 and ARO MURI Award W911NF-08-1-0238.

Keywords and phrases: FARIMA, point process, superposition, multiplexing, long-range dependence, traffic generation, network simulation

typically 40 bytes in size, are added to each packet to manage the file transmission and packet routing. In addition, both hosts can send control packets with no file data, just headers, as part of the transmission management. This means that packet sizes range from 40 bytes to 1500 bytes. Each router has input links and output links; when a packet arrives on an input link, the router reads a field in the header to determine the destination host, and looks in a table to determine the output link over which the packet should be sent to get to the destination.

Each transmission link on the Internet at each point in time can be servicing many ongoing connections. The packet arrival times for transmission on the link are a superposition of the packet arrival times of the individual ongoing connections. As the average number of active connections increases, the packet traffic arrival rate, α packets/sec (p/s), tends to increase. If a packet arrives for transmission and the link is busy transmitting, then the arriving packet is put in a queue. The interface that writes the packet to the link has a speed in bits/sec that determines the service time: the packet size in bits divided by the link speed. The queueing is the major factor in quality-of-service (QoS) for Internet connections; if queueing delay is too large, QoS degrades (Rolls et al., 2005).

1.2. The Critical Role of Packet Arrival Statistical Properties. The statistical properties of the superposed arrival point process are critical because the queueing delay depends heavily on them. In the 1990s, it was discovered in two pioneering articles (Leland et al., 1994; Paxson and Floyd, 1995) that Internet traffic is long-range dependent. The power spectrum as the frequency f goes to zero increases like f^{-2d} for $0 < d < 0.5$. The autocorrelation function as the lag k gets large decreases like k^{2d-1} . These statistical properties make the traffic “bursty”, in the language of network engineering; compared with Poisson arrivals with the same arrival rate, the upper tail of queueing delays is longer, and the amount of traffic that can be put on the link and maintain QoS is less (Duffield, 1996; Erramilli et al., 1996; Heyman and Lakshman, 1996; Park et al., 1997; Ribiero et al., 2006). The development of Internet protocols and devices has been driven in large measure by these statistical properties (Belottia et al., 2008; de Pereira et al., 2002). Modeling the properties has been and remains a critical task for network engineering.

1.3. Past Statistical Foundations and Modeling: Self-Similarity and fGn. Most studies of Internet traffic, including all of the early ones, analyzed packet counts in fixed intervals, a form of time aggregation. In the earliest papers, the traffic was described as self-similar and fractional Gaussian noise (fGn) was put forward as a model (Csabai, 1994; Leland et al., 1994; Paxson and Floyd, 1995; Paxson, 1997; Taqqu et al., 1997b; Willinger et al., 1995, 1996). Using this fGn model, Norros (1994) derived a number of statistical properties, and investigated queueing properties as a part of network engineering study.

fGn is a tractable model and allows mathematical investigations such as the above citation. The problem, however, is that after the earliest papers on Internet traffic, subsequent articles showed that fGn is not an adequate model because traffic is not self-similar across all time scales, and packet counts for small time scales is non-Gaussian.

1.4. *Past Statistical Foundations and Modeling: Multifractal Moment and Time Scaling of m -Blocksums and m -Blockmeans.* The inadequacy of fGn led to intensive study of time scaling properties and the development of new models based on scaling. Work focused on m -blocksums and m -blockmeans, defined as follows. For $u = 1, 2, \dots$, let w_u be a time series. For positive integer m , the m -blocksum process consists of every m -th value of a moving sum of length m :

$$w_v^{(m)} = \sum_{i=1}^m w_{(v-1)m+i}, \quad v = 1, 2, \dots$$

The m -blockmean process is $\bar{w}_v^{(m)} = w_v^{(m)} / m$.

In a very large literature, the statistical properties of these block statistics and how they change with m were studied in many ways (Abry et al., 2002; Ashoura and Le-Ngoc, 2008; Figueiredo et al., 2002; Dang et al., 2003; Erramilli et al., 2002; Gilbert et al., 1999; Feldmann et al., 1998a; Gong et al., 2005; Hannig et al., 2001; Liu and Baras, 2003; Jiang and Dovrolis, 2005; Masugi and Takuma, 2007; Mikosch et al., 2002; Riedi and Vehel, 1997; Riedi et al., 1999; Taqqu et al., 1997a; Roughan and Veitch, 2007; Ribeiro et al., 2005; Stoev et al., 2005; Veitch et al., 2005; Veres and Boda, 2000; Maulik and Resnick, 2003; Karagiannis et al., 2004; Willinger et al., 2002; Yuan et al., 2000; Resnick et al., 2003).

Multifractal wavelet models based on the block statistics were developed (Gao and Rubin, 2001a,b; Riedi et al., 1999; Riedi, 2002; Resnick et al., 2003).

In almost all cases, w_u was taken to be counts in successive small time intervals such as 1 ms or 10 ms. In a few cases, w_u were taken to be interarrival sequences (Gao and Rubin, 2001a; Riedi et al., 1999). The time-aggregation scaling analyses and modeling formed a foundation for intuition about the statistical properties of the arrival time process.

One example of an analysis method is the variance-time plot, which became a very commonly used tool (Erramilli et al., 1996; Fraleigh et al., 2003; Gong et al., 2005; Leland et al., 1994; Riedi et al., 1999): the log of the sample variance of $\bar{w}_v^{(m)}$ is plotted against $\log m$. Another example is autocorrelation-time analysis in which the standard nonparametric estimate of the autocorrelation function of the m -blockmeans is studied as a function of m (Hannig et al., 2001).

Another important method of study was multifractal moment analysis, a study of the moments of normalized values of $w_v^{(m)}$. This was closely associated with

the multifractal wavelet models in the above citations. Observed properties of the moments of the m -blocksums of Internet packet traffic arrivals make the process multifractal (Riedi, 2002). The properties are discussed further in Section 5. The multifractal concept is an enlargement of self-similar processes which have a certain uniformity in the moments that make them monofractal.

Multifractal wavelet models reproduce the statistical properties of traffic arrival times, fixing the shortcoming of fGn. However, there are drawbacks. First, the models are complex and do not have a mathematical tractability that allows derivations of the statistical properties of traffic statistics through mathematical studies. Second, the models are fundamentally nonparametric, requiring $w_v^{(m)}$ when fitting to a interarrival measurements. However, the statistical properties of the packet arrival process change with the traffic rate. Traffic generation by a model for the arrival process for simulation studies must be able to produce traffic at any desired traffic rate. The nonparametric nature of multifractal wavelet models makes them not conducive to the general task of generation at any desired rate.

1.5. GFSD and MFSD Models. This section introduces MFSD models, which have a very simple structure with just three parameters. Coming sections demonstrate that they reproduce the statistical properties of traffic arrival times, are mathematically tractable, provide a new foundation for understanding traffic statistical properties, and can be used to generate packet arrival times at any desired traffic rate.

Let h_u be fractionally differenced white noise, the Hosking discrete analog of fractional Gaussian noise (fGn) (Hosking, 1981),

$$(I - B)^d h_u = \epsilon_u.$$

B is the backward shift operator, $Bh_u = h_{u-1}$; $0 < d < 0.5$ is the fractional-difference power; $(I - B)^d$ is defined by expanding in a power series in B ; and ϵ_u is Gaussian white noise with mean 0, and variance σ_ϵ^2 . We will take

$$\sigma_\epsilon^2 = \frac{(1-d)\Gamma^2(1-d)}{2\Gamma(1-2d)}$$

for purposes stated below. This makes the variance of h_u equal to $(1-d)/2$.

Let s_u be a moving sum of length 2 of h_u ,

$$s_u = h_u + h_{u-1}.$$

s_u can be written in another form,

$$(I - B)^d s_u = \epsilon_u + \epsilon_{u-1},$$

so s_u is a fractional moving-average process (Hosking, 1981). The above value of σ_ϵ^2 makes the variance of s_u equal to 1. Finally, let n_u be Gaussian white noise with variance 1.

A Gaussian fractional sum-difference (GFSD) model for a time series z_u has the form

$$z_u = \sqrt{(1-\theta)s_u} + \sqrt{\theta}n_u,$$

where s_u and n_u are independent processes and $0 \leq \theta \leq 1$. θ is the mixture parameter of the the GFSD. The mean of z_u is 0, and the variance is 1 for notational convenience, and does not limit modeling.

A multifractal fractional sum-difference model (MFSD), t_u , is a stationary discrete time series that is a nonlinear strictly monotone transformation of a GFSD, z_u . Let the cumulative distribution function (cdf) of t_u be $T(t)$, which we suppose is continuous with finite first and second moments. Let $Z(z)$ be the cdf of a Gaussian distribution with mean 0 and variance 1. Then

$$t_u = T^{-1}(Z(z_u)),$$

and

$$z_u = Z^{-1}(T(t_u)).$$

z_u is the Gaussian image of t_u , and t_u is the multifractal image of z_u .

Suppose the marginal distribution of t_u is a Weibull with shape parameter λ . Let the traffic rate be $\alpha = 1/E(t_u)$, measured in packets/sec (p/s). Our parameterization of the Weibull is somewhat different than usual, replacing the usual scale parameter with the rate α , which is more meaningful for packet interarrivals. The cdf for this parameterization is

$$T(t_u) = W(t_u; \lambda, \alpha) = 1 - e^{-\{\alpha\Gamma(1+\lambda^{-1})t_u\}^\lambda}.$$

The transformation to the multifractal image is

$$t_u = \frac{\{-\log(1 - Z(z_u))\}^{1/\lambda}}{\alpha\Gamma(1 + \lambda^{-1})}.$$

t_u is a Weibull MFSD.

Suppose the marginal distribution of t_u is log normal where μ is the mean of $\log(t_u)$ and τ^2 is the variance. The cdf is $T(t_u) = L(t_u; \mu, \tau^2)$. The transformation to the multifractal image is

$$t_u = e^{\tau\sqrt{1-\theta}s_u} e^{\tau\sqrt{\theta}n_u} e^\mu,$$

so the transformation has a simple mathematical form. t_u in this case is a log normal, or multiplicative, MFSD.

Both the Weibull MFSD and the multiplicative MFSD are multifractal, justifying their name; this is demonstrated in Section 5. The Weibull MFSD is the model that is validated here for the packet interarrival process. However, the multiplicative MFSD is simpler and provides a good approximation of the Weibull MFSD; it is used in some cases in coming sections to simplify mathematical derivations.

1.6. *Overview of Results: Validation, Mathematical Study, Foundations, and Traffic Generation.* MFSD and GFSD models were first put forward by Cao et al. (2002). Model parameter estimates were used as summary statistics to demonstrate changing statistical properties of the traffic as the traffic load increases. This article adds three additional sets of results to this initial work.

1.6.1. *Model Validation.* The results include the first extensive validation of the Weibull MFSD as a model for packet interarrivals t_u . Validation is achieved through showing that the statistical properties of Internet traffic statistics derived from the model agree very closely with empirical statistical properties of estimates from live packet trace segments consisting of 715,665,213 measured arrival times on 3 Internet links. The validation study is carried out across a wide range of traffic rates, α , because the statistical properties of the t_u change with α . Such a validation is necessary for the model to be used reliably in the many network engineering studies where accurate traffic models are essential, and to be used reliably for mathematical investigations.

1.6.2. *Mathematical Investigations and Foundations for Traffic Statistical Properties.* The results include mathematical investigations of many traffic statistics, enabled by the mathematical tractability of the model. This leads to closed-form formulas, approximations of these formulas, solutions to equations that are studied numerically, and theorems. An important outcome of the investigations is a new foundation of understanding for the traffic statistics, based on the relative contributions of $\sqrt{1 - \theta} s_u$ and $\sqrt{\theta} n_u$ to the variance of z_u , and how the contributions change with changing factors such as the traffic rate α , and the m of time-aggregation scaling. This constitutes a time-domain foundation for understanding.

1.6.3. *Simple Generation of Traffic.* The results include a method for simple generation of t_u for network engineering traffic studies. Only the traffic rate α needs to be specified. This is achieved by modeling, based on both empirical results and mathematical derivations using the MFSD model, for how the parameters λ , θ , and d change with α . The model for d is a constant. Logit transformations of λ and θ are linear in $\log(\alpha)$.

1.7. *Section Contents.* The contents of the sections are the following: 2. Packet traces that were collected or obtained for the validation study. 3. Marginal distri-

bution of t_u . 4. Power spectrum of z_u . 5. Multifractal properties of t_u . 6. Changes in statistical properties of t_u and z_u with the traffic rate α . 7. Autocorrelation functions of z_u and t_u . 8. Approximations of autocorrelation functions of z_u and t_u . 9. Variance time plot for z_u and t_u . 10. A model is derived for time aggregates of t_u . 11. Summary and discussion.

2. Packet Trace Segments for Model Validation. Validation was carried out by analyzing live packet traces for traffic in both directions of 3 Internet links: Auckland, Leipzig, and Bell. The total number of arrivals is 715,665,213. Bell was the Internet gateway link for a Bell Labs research location with about 500 users. Leipzig was the gateway link for the University of Leipzig campus. Auckland was a link near the edge of the University of Auckland network. All collection used Endace cards (<http://www.endace.com/>) to provide highly accurate, hardware timestamps, which is essential to the modeling. The collected data consist of network and transport headers, and timestamps of packet arrivals, but our analysis used only the packet size field and the timestamp.

The Bell live traces were obtained as a result of one author of this article being a part of the collection operation. The Leipzig live traces were obtained from the Center for Applied Internet Data Analysis (CAIDA) (<http://caida.org/tools>). The Auckland live traces were obtained from the Waikato Internet Traffic Storage (<http://www.wand.net.nz/wits/catalogue.php>).

In coming sections, in the interest of space, we use just Auckland traces in our visual displays and numeric information. However, statistical properties and modeling conclusions were the same for all links. The Auckland traces available for these links were broken into trace segments of 15 min or 1 hr, and each segment analyzed individually. Not all available segments were appropriate for analysis for reasons given below. Table 1 gives information about the analyzed segments; some of the information will be explained later in this section.

2.1. Stationarity. The statistical properties of packet arrivals to a link interface change with the expected number of ongoing connections because superposition of processes changes their statistical properties. The expected number of ongoing connections changes because of calendar variation (day-of-week, holiday, etc.) and diurnal variation in the usage of the measured links. The modeling must account for the changes in the properties. We do this by studying the dependence of the properties on the packet arrival rate α , which is inversely proportional to the expected number of ongoing connections.

To accurately study changes in statistical properties with α , we need each trace segment to have a nearly constant expected rate for the duration of the segment. We insure this, first, by taking segments with small lengths, and then, second, by checking each segment for stationarity by visualization of measures of the packet rate

TABLE 1
Information about analyzed packet trace segments.

Factor	Auckland 15-min	Auckland 1-hr
link speed (megabits/sec)	1000	1000
transmission technology	Ethernet	Ethernet
lab timestamp accuracy (μ sec)	0.030	0.030
collection duration (days)	2.25	2.25
collection date	March 2008	March 2008
number live	96	24
number live-multiplexed	24	0
max utilization live	5.24%	4.95%
min rate live (p/s)	1193	1268
max rate live (p/s)	7674	7712
min rate live-multiplexed (p/s)	2537	NA
max rate live-multiplexed (p/s)	137884	NA

such as the number of packets in 10 sec intervals. We found that 15 min segments worked well, and discarded any segments that showed more than minor nonstationarity. We also used certain 1 hr traces that the visualization showed were close to stationary.

2.2. Packet Rate Above 1000 packets/sec (p/s). Segments whose packets rates are too small, less than about 1000 p/s, are not readily modeled statistically. If there are a small number of ongoing connections, properties of the TCP protocol for individual connections can create cycles in the interarrivals at a number of frequencies of the form of $1/k$ where k is a small integer greater than or equal to 2. These peaks are readily seen in estimates of the power spectrum; their frequencies change across the trace segments, likely due to changes in the Internet application that is dominant. If modeling is needed for very small rates, then a better strategy is to use simulation models that run TCP. This means the MFSD model is not appropriate for packet rates less than about 1000 p/s. We do not model directly trace segments less than this rate, but as discussed below, we can use them in other ways for modeling.

2.3. Modeled Arrivals, Measured Arrivals, and Timestamps. The MFSD model applies to $t_u = a_u - a_{u-1}$ where a_u is the arrival time at the interface queue of the output link. The measured arrival a'_u is the exit time from the queue, and $t'_u = a'_u - a'_{u-1}$ are the measured interarrivals.

If packet u arrives when there is no packet in service, then $a'_u = a_u$. If packet u arrives when a packet is in service, then its transmission begins as soon as packet $u - 1$ has finished; this means that t'_u is equal to the service time of packet $u - 1$. Let p_u be the size of packet u (bits), and let ℓ be the speed (bits/sec) with which the

interface writes a packet to the link. The service time of packet $u-1$ is p_{u-1}/ℓ . This is the smallest possible t'_u when packet $u-1$ has size p_{u-1} . The timestamps, \hat{a}_u , are the a'_u plus measurement error, and the timestamp interarrivals are $\hat{t}_u = \hat{a}_u - \hat{a}_{u-1}$.

2.4. Timestamp Accuracy and Identifying Packets with Queuing Delay. The accuracy of timestamps is critical to the validity of the MFSD modeling. The prediction of timestamp accuracy from laboratory tests of the Endace card used for the Auckland trace segments is $\pm\phi$ where $\phi = 15$ nanosec. This is excellent if it is valid. We will investigate ϕ empirically.

Identifying queued packets is also important for the trace segment selection upon which modeling is based. We selected live trace segments for analysis that have a small percent of delayed packets, less than about 10%, because modeling is for the t_u and not the t'_u . We need trace segments where the t'_u reflect the properties of the t_u . These segments are those with lower packet rates. We determine empirically the percent of queued packets as part of the same method that investigates accuracy.

For all delayed packets u , we have $t'_u = p_{u-1}/\ell$. Measurement errors, however result in timestamps \hat{t}_u of these delayed packets that lie in the interval $p/\ell \pm 2\phi$. Furthermore, we expect that the density of the \hat{t}_u will have a noticeable drop just above $t'_u + 2\phi$. This can lead to a revision in the value of ϕ , and allows identification of packets that experience delay.

An accuracy and delay-identification plot is shown in Figure 1 for the Auckland live trace segment that has the largest bitrate, 33.5 megabits/sec. On the plot, $\hat{t}_u - p_{u-1}/\ell$ is graphed against p_{u-1}/ℓ for u with \hat{t}_u less than 100 nanosec. Because there are 690,239 such u , plotted just a sample of the values. The horizontal lines are drawn at ± 30 nanosec, the laboratory values of $\pm 2\phi$. There is a dense band of points contained within the accuracy limits, and a sharp cutoff in density above the band. This verifies $\phi = 15$ nanosec, and packets within the band can be taken as the queued packets.

2.5. Numerical Multiplexing. To study the changing statistical properties with the packet arrival rate α , we need trace segments with a wide range of observed traffic rates, $\hat{\alpha}$, not just the live 15-min and 1-hr traces whose rates are kept small to ensure a low percent of delayed packets. To achieve larger rate segments, we numerically multiplexed subsets of these 15-min live trace segments to produce numerically-multiplexed 15-min traces with larger rates. This process reflects what happens for the arrival times at the queue, and interarrival times can be arbitrarily small. Figure 2 shows the log packet rates of the 96 live segments and the 24 numerically-multiplexed segments used in our analysis of the Auckland data.

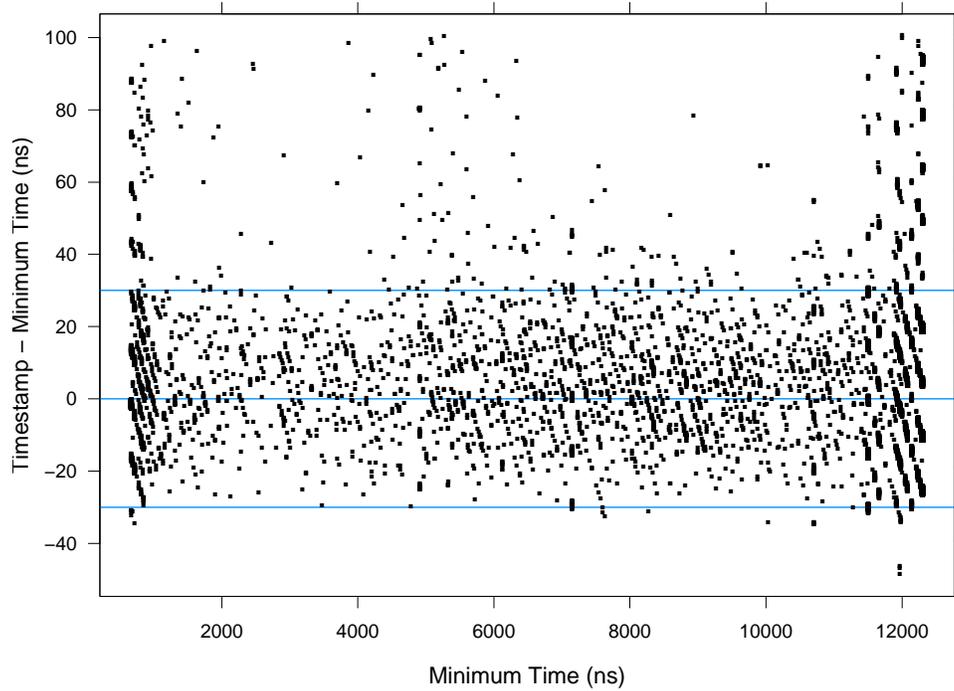


FIG 1. Accuracy and delay-identification plot for one live Auckland 15-minute trace.

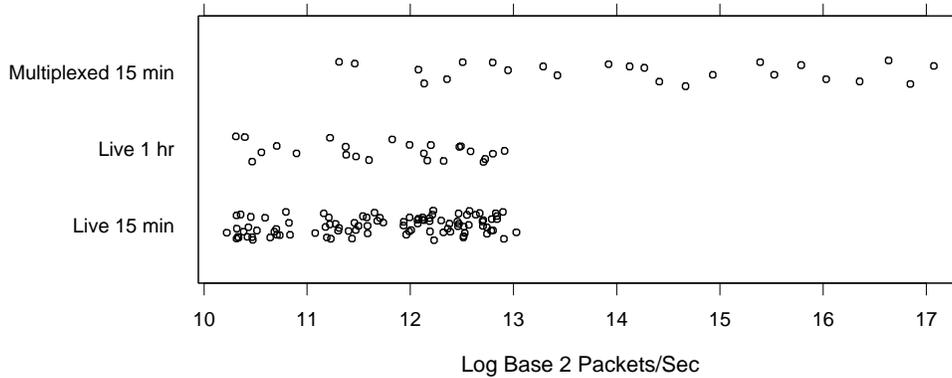


FIG 2. Log base 2 observed packet rate of Auckland trace segments.

2.6. *Visual Displays of 4 Traces.* Data visualization played a critical role in the validation process. There were many types of displays and each was applied to each analyzed trace segment. A number of these display types are shown here for 4 15-min Auckland trace segments to convey results. The packet rates of the 4 segments range from small to large, and are as close to equally spaced on a log

scale as possible. The 2 with the smallest rates are live segments, and the 2 with the largest rates are numerically-multiplexed segments. The packet rates in packets/sec (p/s) are $1771 = 2^{10.79}$, $5634 = 2^{12.46}$, $17928 = 2^{14.13}$, and $66913 = 2^{16.03}$.

3. Multifractal Marginal Distribution: Validation and Properties.

3.1. *Validation of the Marginal Distribution.* The marginal distribution of the interarrival process t_u is well approximated by the Weibull. The parameterization is described in Section 1; α is the packet rate, equal to the inverse of the expected value, and λ is the shape. For each trace segment, λ and α were estimated by the method of moments. $\hat{\alpha}$ is the inverse of the sample mean of the t_u . Using this value, $\hat{\lambda}$ is the value of the Weibull variance that matches the sample variance of the t_u .

The method used to check the Weibull specification is the Weibull quantile plot, illustrated in Figure 3 for the 4 Auckland trace segments described in Section 2.6. In each plot the fourth root of the quantiles of the observed t_u at empirical frequencies 0.00005 to 0.99995 in steps of 0.0001 are plotted against the fourth root of the quantiles of a fitted Weibull using the above estimates. Fourth roots are taken because the resulting transformed distribution is close to symmetric for values of $\hat{\lambda}$ in the range of the trace segments. The vertical lines are drawn at the quantiles with probabilities 0.01, 0.05, 0.25, 0.75, and 0.95, and 0.99. The oblique line has slope 1 and intercept 0.

If the observed t_u are well approximated by a Weibull, then the pattern of the points on the plot should follow the oblique line. In Figure 3, and for almost all other analyzed trace segments, the Weibull provides an excellent fit, taking sampling variability and artifacts into account. There are small departures, atoms in the live empirical distributions, in the top 2 panels. These are artifacts resulting from up to about 10% of the measured interarrivals not being the same as the modeled interarrivals. This is discussed in Section 2; queueing on the link input interface results in noticeable atoms in the measured interarrivals equal to p/ℓ for commonly occurring packet sizes p . This is nearly eliminated in the bottom two panels due to the numerical multiplexing.

3.2. *The Change in λ with α .* Let $\hat{\lambda}_k$, $k = 1, \dots, 144$, be the estimates of the shape λ for the 144 Auckland trace segments, and let $\hat{\alpha}_k$ be the estimates of the packet rate α . Figure 4 graphs $\hat{\lambda}_k$ against $\log_2(\hat{\alpha}_k)$ where \log_2 is log base 2. The smallest values of $\hat{\lambda}_k$ are close to 0.6; they tend to 1 as $\log_2(\hat{\alpha}_k)$ increases, which means the marginal distribution tends to exponential.

Section 6 presents a derivation of λ as a function of α using the MFSD model. Equations are solved that yield numeric values, leading to a model $\lambda(\alpha)$ for the dependence of λ on α . The theoretical model agrees with the empirical pattern in Figure 4. This dependence of λ on α is a critical aspect of the statistical properties

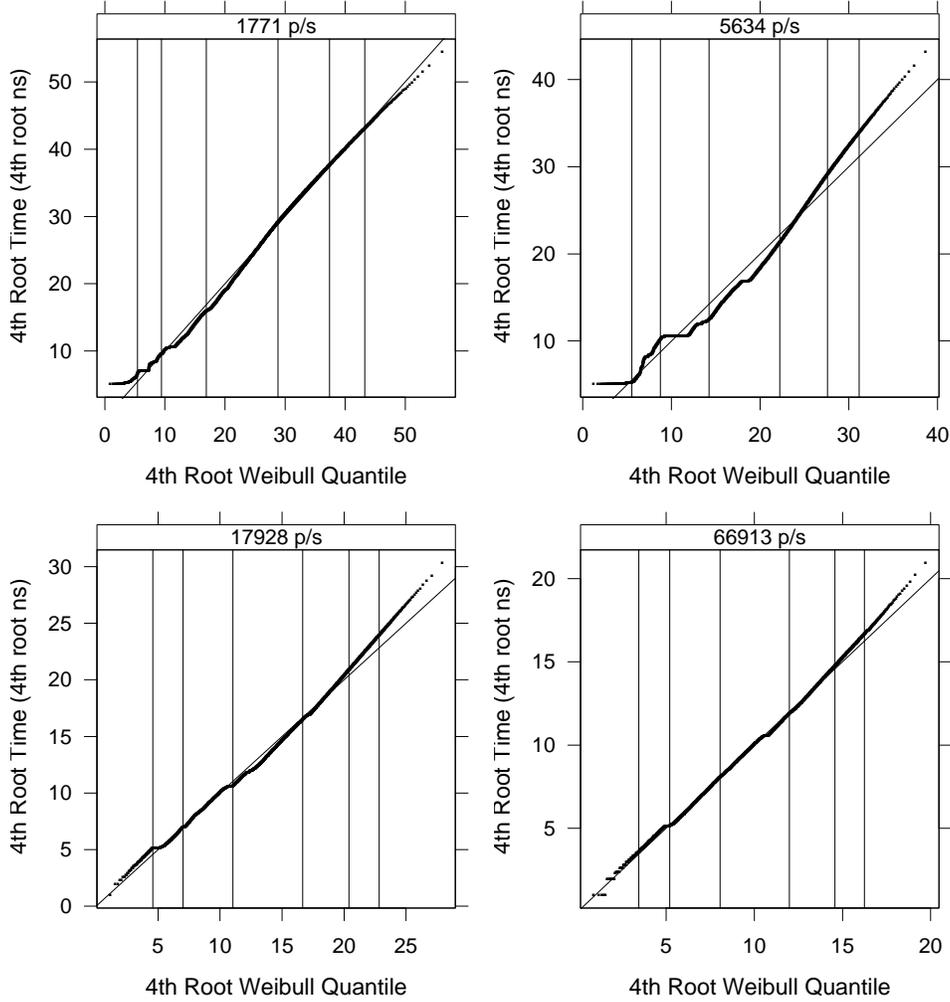


FIG 3. Weibull quantile plots of 4 trace segments.

of the packet arrival process, so we switch notation from λ to $\lambda(\alpha)$ in coming sections.

4. Gaussian Power Spectra: Validation and Properties. The four time series considered in the GFSD, which are defined in Section 1, are h_u , s_u , n_u , and z_u . This section presents formulas for their power spectra that provide insight about statistical properties. Validation study is also carried out for the observed z_u of each trace segment by comparing nonparametric estimates of the power spectrum with that of a GFSD model fitted to the z_u .

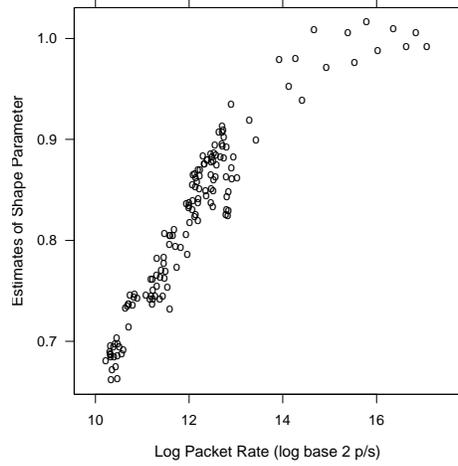


FIG 4. $\hat{\lambda}$ is plotted against $\log_2(\hat{\alpha})$ for 96 live 15-min trace segments (\circ), 24 live 1-hr segments (\circ), and 24 numerically-multiplexed segments (\circ).

4.1. *Formulas and Statistical Properties.* One outcome of the simplicity of the GFSD is simple formulas for their power spectra. Let $0 < f \leq 0.5$ denote frequency in units of cycles/interarrival. The power spectra are

$$\begin{aligned} p_h(f) &= \frac{(1-d)\Gamma^2(1-d)}{2\Gamma(1-2d)\{2\sin(\pi f)\}^{2d}} \\ p_s(f) &= 4\cos^2(\pi f)p_h(f) \\ p_n(f) &= 1 \\ p_z(f) &= (1-\theta)p_s(f) + \theta. \end{aligned}$$

d is the fractional difference exponent and θ is the mixture parameter. $p_h(f)$, $p_s(f)$, and $p_z(f)$ decrease strictly monotonically as f increases, and all go to infinity to order f^{-2d} at the origin, a signature property of the long-range dependence amply observed empirically in many previous studies.

There are an infinite number of ways of decomposing z_u into a long-range dependent component plus a white noise component. The decomposition of the GFSD,

$$z_u = \sqrt{1-\theta}s_u + a\sqrt{\theta}n_u,$$

is the one that maximizes the variance of the white noise because $p_s(0.5) = 0$. This means $p_z(0.5) = \theta$, which will be used below in the estimation of θ .

Figure 5 graphs $\ell_z(f) = 10 \log_{10}\{p_z(f)\}$ against f where \log_{10} is log base 10. In visual displays of the power spectra, we switch to this decibel scale because

it shows features more effectively. Three spectra are plotted. For each, $d = 0.31$ because estimates of d change little, apart from statistical variability, across the trace segments, and 0.31 is a central value. The values of θ are different for the three spectra: 0.6 (—), 0.8 (—), 0.975 (—). These values reflect the range of the estimates of θ across trace segments. The spectra were evaluated at values of f from 2^{-16} to $1/2$. The vertical line on each panel is drawn at frequency $f_0 = 0.129$ for reasons we explain next.

For fixed f and d , $p_z(f)$ is linear in θ with derivative $1 - p_s(f)$. Let f_0 be the frequency where $1 - p_s(f_0) = 0$, which means $p_z(f_0)$ and $\ell_z(f_0)$ do not change with θ . f_0 depends only on d , and for $d = 0.31$, $f_0 = 0.129$ cycles/interarrival, which has a period of 7.75 interarrivals. This is the value at which the vertical line is drawn in Figure 5, and we can see that the spectra do not change with θ at this frequency. It is easy to see that $p_z(f)$ and $\ell_z(f)$ decrease with θ for $f > f_0$, and increase for $f < f_0$. This is also demonstrated in Figure 5.

4.2. *Estimation of Parameters d and θ .* To carry out estimation and model checking for the GFSD model for each trace segment, the observed t_u for each segment were transformed to observed z_u by the function $z_u = Z^{-1}\{\hat{T}(t_u)\}$, where \hat{T} is the empirical cumulative distribution function of the t_u , and Z is the normal cumulative distribution function with mean 0 and variance 1. Let n be the number of t_u in the segment. Let $r(u)$ be the rank of t_u . Then $z_u = Z^{-1}\{(r(u) - 0.5)/n\}$. The reason for using the empirical function, rather than a Weibull distribution function fitted to the t_u , was to have a portion of the model checking methods for z_u not depend on the validity of the specification of the marginal distribution of t_u . In

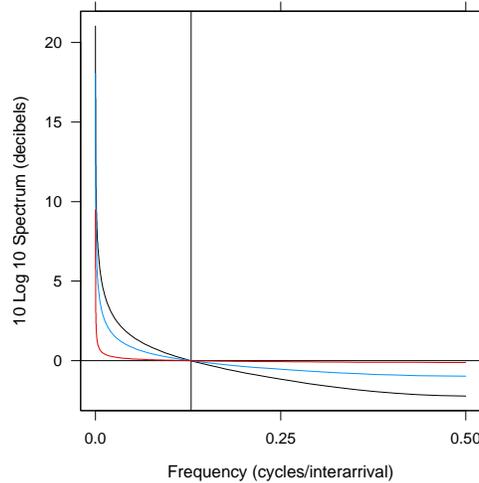


FIG 5. Log power spectra $\ell_z(f)$ for $d = 0.31$ and 3 values of θ .

Section 3, model checking of the marginal of t_u does not depend on the validity of the specifications for z_u . In a number of other sections, model checking depends jointly on specifications for the multifractal and Gaussian images. Estimation of the parameters d and θ of the GFSD are based on the periodogram of the observed z_u . In addition, the periodogram and the estimate of the power spectrum using the estimated parameters are a part of our model checking for validation of the GFSD.

To enable subsequent study of the m -blockmeans and m -blocksums with exactly the same observations, we use just the first 2^b observations of each trace segment where b is the greatest integer in log base 2 of the number of z_u in the segment. For the 144 Auckland trace segments, the minimum value of b is 20 and the maximum is 26. The periodogram is computed at the Fourier frequencies $f_i = i/2^b$ for $i = 1, 2, 3, \dots, 2^{b-1}$. These frequencies are divided into 2^{15} non-overlapping blocks of equal length, so each has 2^{b-16} values. For $j = 1, \dots, 2^{15}$, let \bar{f}_j be the mean of the frequencies in block j , and let $\bar{I}(\bar{f}_j)$ be the mean of the periodogram values in the block. Estimation and model checking proceed with \bar{f}_j and $\bar{I}(\bar{f}_j)$.

Our parameter estimation method for each trace segment is designed to be robust to minor departures of the patterns in the $\bar{I}(\bar{f}_j)$ from the general form of the GFSD power spectrum. Some departures can adversely affect the estimation of d (Hurvich et al., 2002). For example, minor low-frequency trends can remain because the detrending methods described in Section 2 cannot entirely remove the diurnal variation in the packet rate α .

The estimate $\hat{\theta}$ of θ is taken to be the mean of the $\bar{I}(\bar{f}_j)$ for $f_j \geq 0.48$, since $p_z(0.5) = \theta$. This insures that the estimated power spectrum fits the pattern of the $\bar{I}(\bar{f}_j)$ for high frequencies. d is estimated from another frequency band: $0.01 \leq \bar{f}_j \leq 0.06$. \hat{d} is the estimate arising from a nonlinear least squares fit of values of $10 \log_{10}(p_z(f_j))$ with $\theta = \hat{\theta}$ are fitted to the values of $10 \log_{10}(\bar{I}(\bar{f}_j))$ for f_j in the band. This is a variation of the method of (Geweke and Porter-Hudak, 1983) where the frequency band is $0 < f < a$ for a small a . The averaging of the periodogram before taking the log in the least-squares fitting falls in the category of an ATS method (Cleveland et al., 1993); averaging before moving to a log scale results in efficient least-squares estimation.

4.3. *The Change in θ and d with α .* For the 144 Auckland trace segments, the left panel of Figure 6 graphs the estimates $\hat{\theta}_k, k = 1, \dots, 144$ of θ against the log estimates of the packet rates, $\log_2(\hat{\alpha}_k)$, where \log_2 is log base 2. The right panel graphs \hat{d}_k against $\log(\hat{\alpha}_k)$. The estimates $\hat{\alpha}_k$ use the method of moments described in Section 3. The smallest values of $\hat{\theta}_k$ are close to 0.6; they tend to 1 as $\log_2(\hat{\alpha}_k)$ increases, which means that z_u tends to white noise. Except for two large outliers, values of \hat{d}_k vary from about 0.28 to 0.35, a narrow range. The median, shown by the horizontal line, is 0.31. This suggests that d does not change appreciably with α

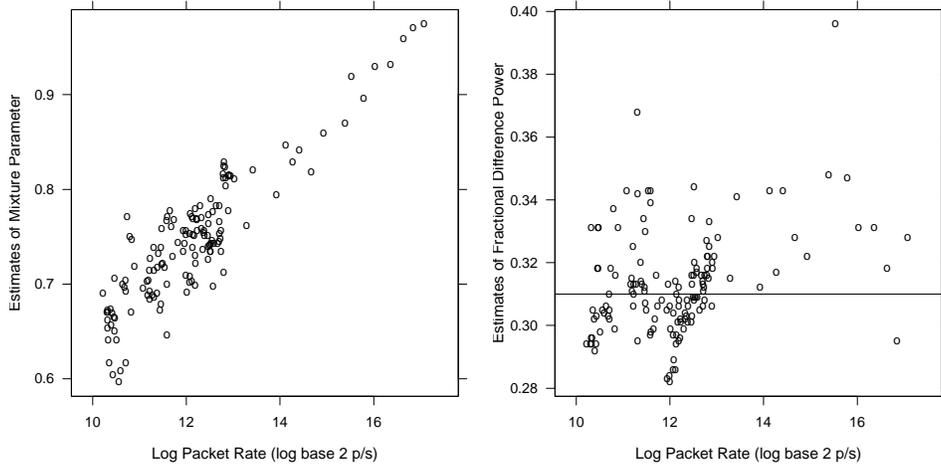


FIG 6. Estimates $\hat{\theta}_k$ (left panel) of the mixture parameter and estimates \hat{d}_k (right panel) of the fractional difference power vs. observed log traffic rates $\log_2(\hat{\alpha}_k)$ for the 144 Auckland traces. The horizontal line in the right panel shows the median, 0.31, of the \hat{d}_k .

so that a fixed d of 0.31 is reasonable in our mathematical study of traffic statistics based on the MFSD model.

Section 6 presents a derivation of θ as a function of α using the MFSD model. Equations are solved that yield numeric values, leading to a model $\theta(\alpha)$ for the dependence of λ on α . The theoretical model agrees with the empirical pattern in Figure 6. This dependence of θ on α is a critical aspect of the statistical properties of the packet arrival process, so we switch to the notation $\theta(\alpha)$ in coming sections.

4.4. *Model Validation: Properties of the Power Spectrum.* The validity of the GFSD model — its ability to account for the statistical time-series properties of z_u — was explored by studying power spectra, one description of the properties. Other descriptions are studied in later sections.

For power spectra, a visual diagnostic method for each trace segment compares the following: (1) $10 \log_{10}(\bar{I}(\bar{f}_j))$, which is a (noisy) nonparametric estimate of the log power spectrum; (2) $10 \log_{10}(\hat{p}_z(\bar{f}_j))$ with $\theta = \hat{\theta}_i$ and $d = \hat{d}_i$, which is the GFSD model estimate of the power spectrum. Figure 7 show the results for the 4 Auckland traces described in Section 2.6. Each panel of the top row graphs $10 \log_{10}(\bar{I}(\bar{f}_j))$ (\bullet) and $10 \log_{10}(\hat{p}_z(\bar{f}_j))$ ($—$) against \bar{f}_j . The bottom row is similar, except that values are graphed against $10 \log_{10}(\bar{f}_j)$. The fits are quite good. The GFSD model estimates do a good job of fitting the patterns of the nonparametric estimates. This was the case for almost all of the packet trace segments of our validation study.

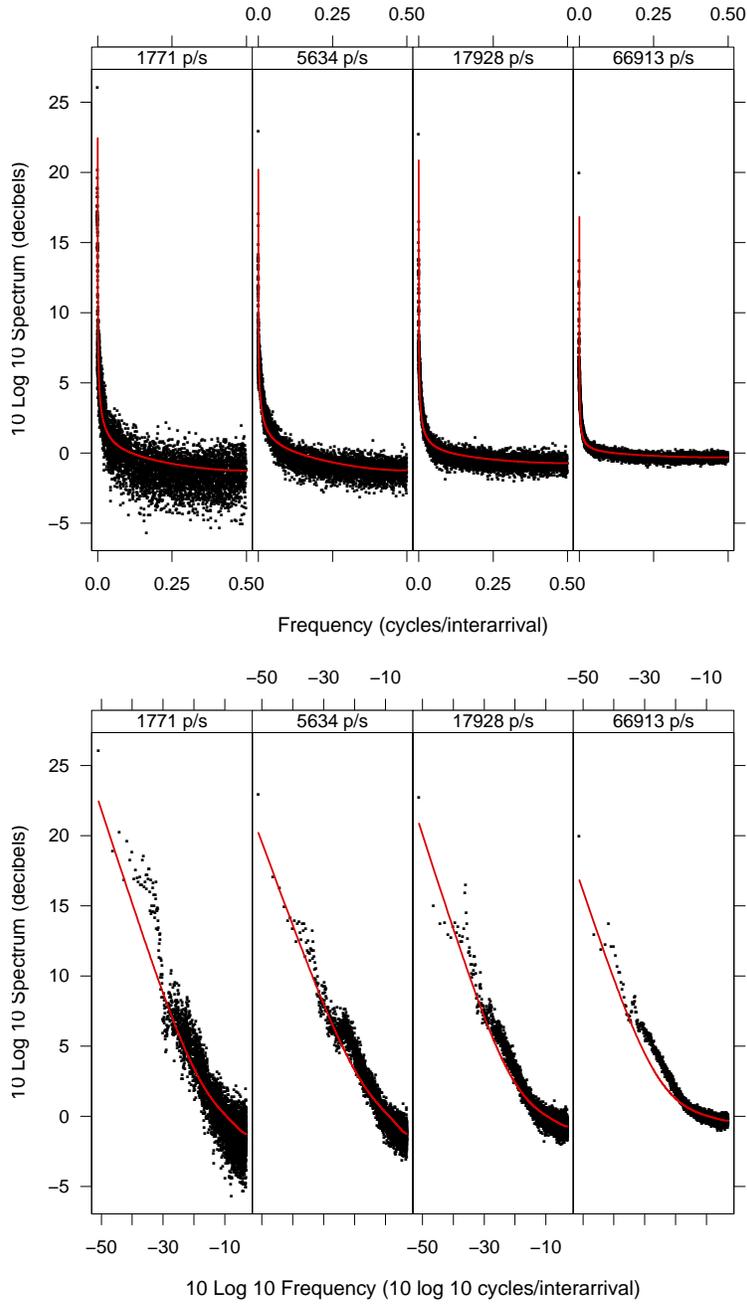


FIG 7. $10 \log_{10}(\bar{I}(\bar{f}_j))$ (\bullet) and $10 \log_{10}(\hat{p}_z(\bar{f}_j))$ (---) for 4 trace segments.

5. Multifractal m -Blocksum Moment-Statistics: Validation. This section addresses the nonlinearity of t_u through a moment-statistic study of the m -blocksum process $t_v^{(m)}$ described in Section 1. Normalized estimates of $E\{(t_v^{(m)})^q\}$ are studied as a function of q and m , which is a time scaling study for each q . This multifractal analysis is a standard in the Internet traffic literature. For each trace segment we compared nonparametric moment-statistics of the t_u with the theoretical moment-statistics from both the Weibull and multiplicative MFSD models fitted to the t_u . This provides an important look at nonlinear properties to aid validation, which very much justifies the analysis. However, it does not provide the foundational insights that arise from analyses in other sections.

For each trace segment, we estimated moments using the first $n = 2^b$ observations of t_u where b is the largest integer in the log base 2 of the number of interarrivals, the same data selection method used in Section 4. Let $t. = \sum_{u=1}^n t_u$. The nonparametric q -th moment estimate for the m -blocksum is

$$(1) \quad \hat{S}_q^{(m)} = \sum_{v=1}^{2^{b-r}} \left(\frac{t_v^{(m)}}{t.} \right)^q.$$

Values of m were $m_r = 2^r$ for $r = 0, \dots, b - 1$, and the moments were $q = -10, -5, -2, 2, 5, 10$.

The Weibull MFSD has 4 parameters. Two are for the Weibull marginal of t_u : the shape λ and the packet rate α . Two are for the associated Gaussian image z_u : the fractional difference coefficient d and the mixture parameter θ . The fitted Weibull MFSD for each trace segment is the MFSD with parameter values equal to the estimates described in Sections 3 and 4: $\hat{\alpha}$, $\hat{\lambda}$, \hat{d} , and $\hat{\theta}$. The multiplicative MFSD has 4 parameters. Two are for the lognormal marginal of t_u : the mean μ and variance τ of $\log(t_u)$. Their estimates $\hat{\mu}$ and $\hat{\tau}$ are the values for which the first and second moments of the log normal match the second moments of the Weibull with parameters $\hat{\alpha}$ and $\hat{\lambda}$. Two are for the associated Gaussian image z_u : d and θ . Their estimates are also those of Section 4: \hat{d} , and $\hat{\theta}$.

We are unable to mathematically derive MFSD moment-statistics for the Weibull and multiplicative MFSD models, so simulation “derivations” were carried out for each trace segment. Each run for a trace segment consisted of generation of 2^b values of t_u from the fitted model, the same number used for the nonparametric moment-statistics. Moment statistics for the run are computed using Equation 1, and final estimates, $\hat{S}_q^{(m)}$, are means across 100 runs.

Figure 8 is a moment-statistic m -plot for one of the four trace segments described in Section 2, the one with packet rate $\hat{\alpha} = 17928$ p/s. The other 3 segments of the section are not shown in the interest of space. Each panel plots $\log_2\{\hat{S}_q^{(m)}\}$ against $\log_2\{m_r\}$ for one of three cases: nonparametric, Weibull MFSD, and mul-

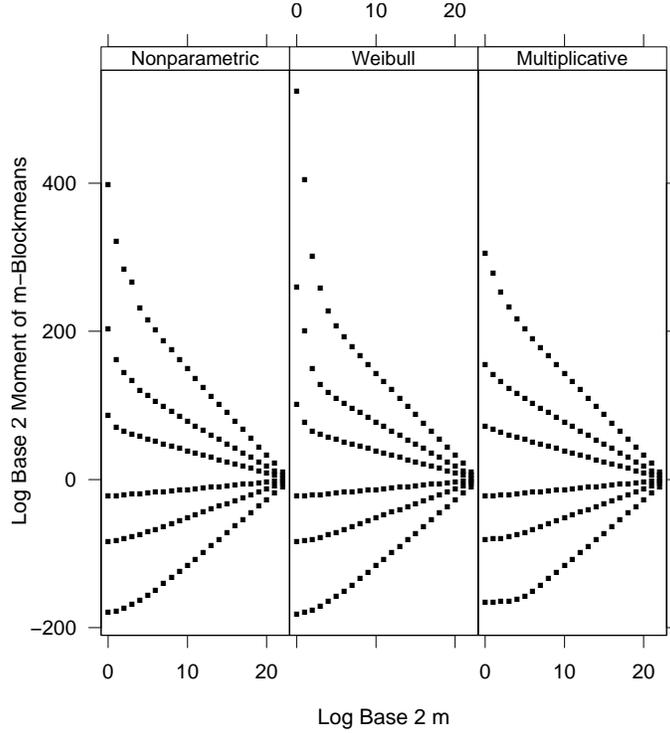


FIG 8. Moment-statistic m -plot for one Auckland trace..

tiplicative MFSD. Successive values for each q are connected by line segments. The resulting “curves” increase with q .

Figure 9 is a moment-statistic q -slopeplot. Let

$$\kappa_{r,r+1}(q) = \frac{\log_2\{\hat{S}_q^{(m_r)}\} - \log_2\{\hat{S}_q^{(m_{r+1})}\}}{\log_2\{m_r\} - \log_2\{m_{r+1}\}}$$

for $r = 0, \dots, b - 2$. In Figure 9, $\kappa_{r,r+1}(q)$ is plotted against q for the first 8 slopes. Each of the 3 panels in Figure 8 results in one row of panels in Figure 9. The values of $[r, r + 1]$ are shown in the strip label of each panel. The line on each panel of the figure goes through the first two points plot to help judge linearity.

The most important aspect of Figures 8 and 9 is that the patterns for the nonparametric, Weibull MFSD, and multiplicative MFSD moment-statistics are very close, and agree with patterns of nonparametric moment-statistics seen in previous publications (Feldmann et al., 1998a; Gao and Rubin, 2001a; Riedi et al., 1999). This provides a striking validation for this model checking analysis. The nonlinearity of $\kappa_{r,r+1}(q)$ as a function q shown in the panels Figure 9 are an indicator of multifrac-

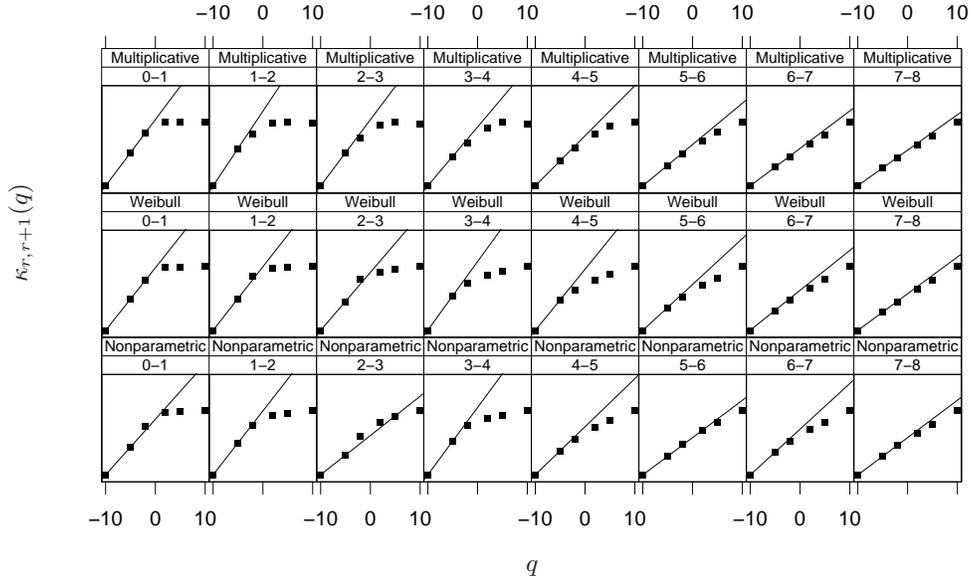


FIG 9. Moment-statistic q -slopeplot for one Auckland trace.

tal behavior. The patterns are concave for $[0, 1]$ and tend toward linear as $[r, r + 1]$ increases. Patterns for $[8, 9]$ and above, not shown, are very close to linear.

6. Multifractal Traffic Generation: Modeling Changes in λ and θ with α .

The statistical properties of t_u change with an increase in the traffic rate α because the expected number of active connections increases with α , which means more multiplexing, or superposition, of the packets from different connections. The parameters of the Weibull MFSD — the Weibull shape parameter $\lambda(\alpha)$ and the Gaussian image mixture parameter θ — reflect this change. Sections 3 and 4 show that each tends to 1 with α ; the fractional difference coefficient d does not change appreciably with α and is taken to be 0.31.

We can also study the change in $\lambda(\alpha)$ and $\theta(\alpha)$ with α theoretically using the Weibull MFSD model. We do this in two ways. The first is a derivation by simulation in which traffic is generated using the Weibull MFSD model. The second is a heuristic mathematical derivation whose detail is described in the Appendix (Section 13). For both, we fix $d = 0.31$ and use initial values $\lambda_0 = 0.70$ and $\theta_0 = 0.55$ at the traffic rate $\alpha_0 = 2^{10.22}$ p/s, the smallest rate for the Auckland trace segments. The initial values were chosen so that the derivations provide the best fit to the estimates $\hat{\lambda}$ and $\hat{\theta}$ of Sections 3 and 4 as functions of the packet rate estimates $\hat{\alpha}$.

For the simulation, we generated 2 Weibull MFSD series, each with parameters

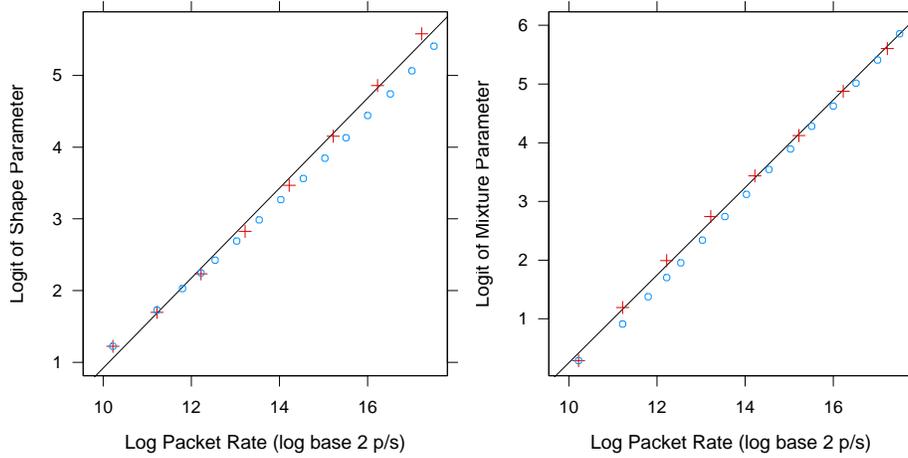


FIG 10. Logit transformations of λ (left) and θ (right), derived by simulation (+) and derived mathematically (\circ), are plotted against $\log_2(\alpha)$.

λ_0 and θ_0 for the rate α_0 . The two MFSD series were then numerically multiplexed, forming a series with rate $2^{11.22}$ p/s. $\lambda(\alpha)$ and $\theta(\alpha)$ were estimated using the methods employed in Sections 3 and 4 for the live and numerically multiplexed data, but with d fixed at 0.31. Then two series were generated at rate $2^{11.22}$ p/s using the estimated parameters, these two series were multiplexed, and then the parameters again estimated. This process continued up to rate $2^{17.22}$ p/s. The result is 8 values of $\lambda(\alpha)$ and $\theta(\alpha)$ including the initial values, and 8 associated values of α .

For the mathematical derivation, the process proceeds in a similar way, but with a different multiplexing method. r Weibull MFSD series with rates α_0 and parameters λ_0 and θ_0 were assumed to be multiplexed. Then values of $\lambda(\alpha)$ and $\theta(\alpha)$ for the multiplexed series were derived. The values of r were 2, 3, 4, 5, 7, 10, 14, 20, 28, 39, 55, 78, 110, 155. The rates for the derived parameters are from $2^{10.22}$ p/s to $2^{17.5}$ p/s. The result is 15 values of $\lambda(\alpha)$ and $\theta(\alpha)$ including the initial values, and 15 associated values of α .

Figure 10 graphs logit transformations of the derived values of the parameters, $\text{logit}_2(\lambda(\alpha)) = \log_2\{\lambda(\alpha)/(1 - \lambda(\alpha))\}$ and $\text{logit}_2(\theta(\alpha)) = \log_2\{\theta(\alpha)/(1 - \theta(\alpha))\}$, against $\log_2(\alpha)$. Each panel shows the simulation derived values (+) and the mathematically derived values (\circ) for one parameter. The results of the two derivations for each parameter as a function of rate are very close. This is an important validation of the mathematically derived values because certain assumptions are made that are not true for an Weibull MFSD model, but that are believed not to affect the results. The logit transformation results in a nearly linear dependence on $\log_2(\alpha)$. The line on each panel is the least squares fit to the simulated values. The

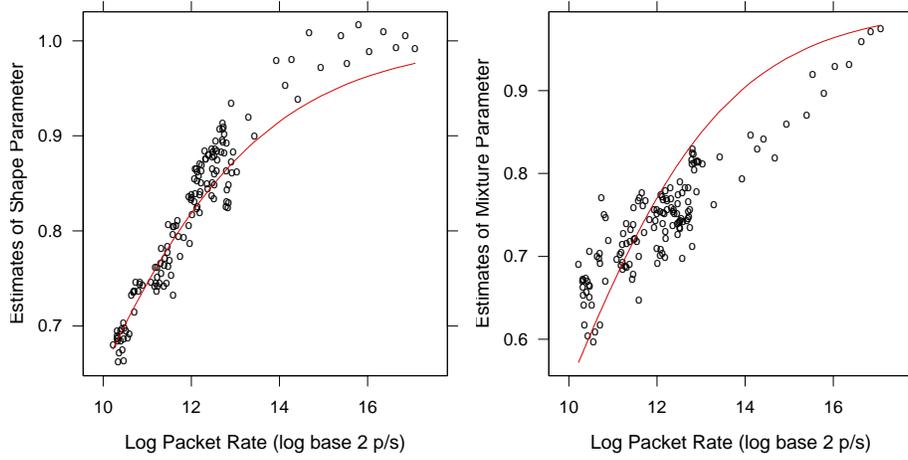


FIG 11. $\hat{\lambda}$ (left panel \circ) and $\hat{\theta}$ (right panel \circ) are plotted against $\log_2(\hat{\alpha})$. Derived equations $\lambda(\alpha)$ (left panel —) and $\theta(\alpha)$ (right panel —) are plotted against $\log_2(\alpha)$.

equations are

$$(2) \quad \text{logit}(\lambda(\alpha)) = -5.36 + 0.63 \log_2(\alpha)$$

$$(3) \quad \text{logit}(\theta(\alpha)) = -7.21 + 0.75 \log_2(\alpha).$$

The equations on the scales of the parameters are

$$(4) \quad \lambda(\alpha) = \frac{2^{-5.36} \alpha^{0.63}}{1 + 2^{-5.36} \alpha^{0.63}}$$

$$(5) \quad \theta(\alpha) = \frac{2^{-7.21} \alpha^{0.75}}{1 + 2^{-7.21} \alpha^{0.75}}$$

Equations 4 and 5 can be used for generating traffic using the Weibull MFSD with a specification of the packet traffic rate α only.

Figure 11 graphs $\hat{\lambda}$ and $\hat{\theta}$ against $\hat{\alpha}$ for the 144 Auckland traces. The curves are an evaluation of Equations 4 and 5 plotted against $\log_2(\alpha)$. There is substantial variability in the estimates, both the ordinates and the abscissas of the plot. The curves do a reasonable job of fitting the patterns of the estimates considering this variability.

7. Gaussian Autocorrelations: Validation and Properties. The four time series considered in the GFSD, which are defined in Section 1, are h_u , s_u , n_u , and z_u . This section presents formulas for their autocorrelations, which sets the stage for Section ?? where approximations of the autocorrelations are derived for both z_u and t_u that provide important insight about statistical properties.

This section also describes results of the validation study carried out for the observed z_u of each trace segment by comparing nonparametric estimates of the autocorrelation function with that of a GFSD model fitted to the z_u . This parallels the analysis of Section 4 that used the power spectrum for model checking. Mathematically, the autocorrelation function is equivalent to the power spectrum in that each is a Fourier transform of the other, but both are used for validation since a small consistent departure across lags or frequencies of one can translate to a large departure locally at certain frequencies or lags of the other.

The autocorrelation function for n_u is $\rho_n(k) = 0$. The 3 other series, which are long-range dependent, have formulas that are easily derived from results of Hosking (1981). The autocorrelation at lag k for h_u , s_u and z_u , respectively, are

$$\begin{aligned}\rho_h(k) &= \frac{\Gamma(1-d)}{\Gamma(d)} \frac{\Gamma(k+d)}{\Gamma(k-d+1)} = \prod_{i=1}^k \frac{(d+i-1)}{i-d} \\ \rho_s(k) &= \rho_h(k) \frac{2(1-d)k^2 - (1-d)^2}{k^2 - (1-d)^2} \\ \rho_z(k) &= (1 - \theta(\alpha))\rho_s(k).\end{aligned}$$

$\rho_z(k) > 0$ and $\rho_z(k)$ goes to 0 with k to order k^{2d-1} , a signature property of the long-range dependence amply observed empirically in many studies.

The validation process for each trace segment begins, as described in Section 4, with a transformation to the observed Gaussian image z_u from the observed multifractal image t_u for $i = 1, \dots, n$. The nonparametric estimate of autocorrelation at lag k for a segment is

$$n^{-1} \sum_{u=1}^{n-k} z_{u+k} z_u.$$

The fitted GFSD autocorrelations are an evaluation of Equation 6 using the estimates \hat{d} and $\hat{\theta}$ from Section 4.

Figure 12 graphs the nonparametric estimates (\bullet) and the fitted GFSD estimates ($—$) against \sqrt{k} for the 4 Auckland trace segments described in Section 2. The square-root lag is used because it allows better assessment of the autocorrelations for small lags. The fitted GFSD estimates provide an excellent fit to the nonparametric estimates. This is the case for for almost all trace segments.

8. Self-Similarity: h_u and s_u .

8.1. Introduction.

8.2. *Approximating the Autocorrelation Functions.* The fGn process, which is the stationary increment of self-similar fBm, has an autocorrelation function

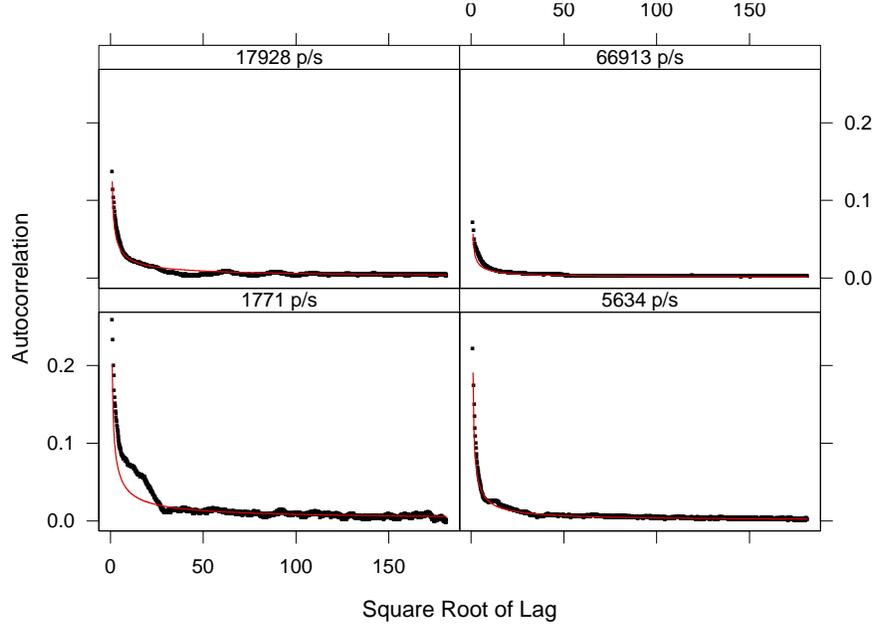


FIG 12. Nonparametric estimates (\bullet) and GFSD fitted model estimates ($—$) of autocorrelation are plotted against square root lag for 4 Auckland trace segments.

approaching $d(2d + 1)k^{2d-1}$ for $k \rightarrow \infty$, a very simple mathematical form that allows much insight and tractable mathematics (Norros, 1994). It has the attractive property that the log of the autocorrelation is linear in the log of the lag. Hosking proposed h_u as a discrete analog of continuous fGn.

A first question is whether the autocorrelation function of h_u ,

$$\rho_h(k) = \frac{\Gamma(1-d)}{\Gamma(d)} \frac{\Gamma(k+d)}{\Gamma(k-d+1)},$$

is well approximated by a constant times k^{2d-1} . From Stirling's formula,

$$\lim_{k \rightarrow \infty} \frac{\Gamma(k+d)/\Gamma(k-d+1)}{k^{2d-1}} = 1,$$

so we approximate by

$$\check{\rho}_h(k) = \frac{\Gamma(1-d)}{\Gamma(d)} k^{2d-1}.$$

This is not the only possibility; for example, we could attempt an approximation in which the constant is chosen so that $\check{\rho}_h(1) = \rho_h(1)$.

We saw in Section 4 that the estimates, \hat{d} , of the fractional difference parameter d , varied by a small amount, and were centered on a median of 0.31. Thus it is

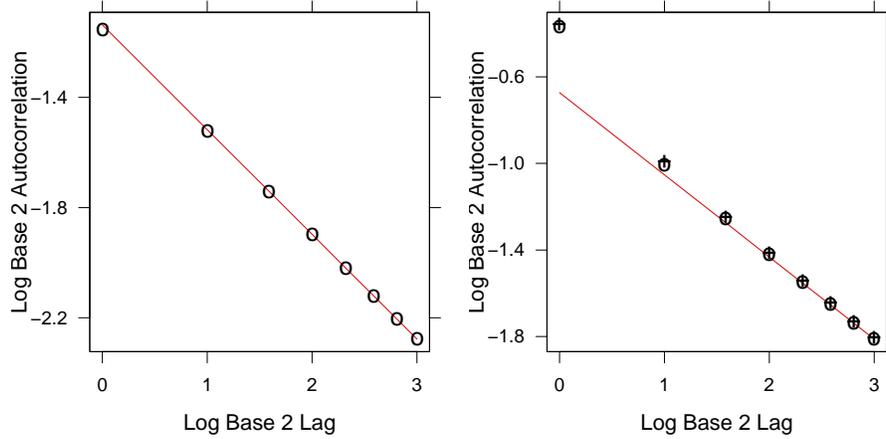


FIG 13. Log base 2 autocorrelation vs. log base 2 lag. Left: $\rho_h(k)$ (\circ), $\check{\rho}_h(k)$ ($-$). Right: $\rho_s(k)$ (\circ), $\dot{\rho}_s(k)$ ($+$), $\check{\rho}_s(k)$ ($-$).

entirely reasonable to take $d = 0.31$ for the GFSD model for z_u . The left panel of Figure 13 plots $\log_2(\check{\rho}_h(k))$ ($-$) and $\log_2(\rho_h(k))$ (\circ) against $\log_2(k)$ for $k = 1, \dots, 8$. As we can see, the very simple $\check{\rho}_h(k)$ is an excellent approximation for these 8 lags. The largest discrepancy is at $k = 1$: $\rho_h(1)/\check{\rho}_h(1) = 1.012$. For $k = 2, \dots, 8$, the discrepancy decreases, and for $k > 3$ is negligible. The approximation is excellent; it could be used for calculations for most purposes, although it is typically not necessary.

For s_u , we consider two approximations of $\rho_s(k)$. Since $s_u = h_u + h_{u-1}$,

$$\rho_s(k) = \frac{1-d}{2} \{ \rho_h(k-1) + 2\rho_h(k) + \rho_h(k+1) \}.$$

The first approximation uses $\check{\rho}_h(k)$ in place of $\rho_h(k)$ in this last equation:

$$\dot{\rho}_s(k) = \frac{(1-d)\Gamma(1-d)}{2\Gamma(d)} \{ (1-1/k)^{2d-1} + 2 + (1+1/k)^{2d-1} \} k^{2d-1}$$

The second simplifies by replacing each of the two terms $(1-1/k)^{2d-1}$ and $(1+1/k)^{2d-1}$ by 1:

$$\check{\rho}_s(k) = \frac{2\Gamma(2-d)}{\Gamma(d)} k^{2d-1}.$$

Both approximations are exact in the limit,

$$\lim_{k \rightarrow \infty} \frac{\dot{\rho}_s(k)}{\rho_s(k)} = \lim_{k \rightarrow \infty} \frac{\check{\rho}_s(k)}{\rho_s(k)} = 1.$$

$\log_2(\rho_s(k))$ (o), $\log_2(\dot{\rho}_s(k))$ (+), and $\log_2(\ddot{\rho}_s(k))$ (−) are plotted against $\log_2 k$ for $d = 0.31$ and $k = 1, \dots, 8$ in the right panel of Figure 13. $\rho_s(k)$ and $\dot{\rho}_s(k)$ are in very close agreement for all lags; the latter certainly can be used for faster computation. For the second approximation, we have $\rho_s(k)/\ddot{\rho}_s(k)$ for $k = 1, 2, 3, 4$ are 1.235, 1.034, 1.014, and 1.007. For $k > 4$, the differences are negligible, and $\ddot{\rho}_s(k)$ could be used for computation for these lags. In addition, for heuristic reasoning about s_u , the closeness of the approximation means we can think of s_u as a near-fGn process whose statistics are very close to those of h_u .

8.3. *m-Blocksums.* In this Section we show that $s_v^{(m)}/m^{d+0.5}$ has the same distribution as $s_v^{(3)}/3^{d+0.5}$. The same holds for $h_v^{(m)}/m^{d+0.5}$.

8.3.1. *Variance.* To begin, we have

$$(6) \quad V(s_v^{(m)}) = m + 2 \sum_{k=1}^m (m-k) \rho_s(k),$$

We saw that $\ddot{\rho}_s(k) = \{2\Gamma(2-d)/\Gamma(d)\}k^{2d-1}$ provides a good approximation of $\rho_s(k)$. We use this to approximate $V(s_v^{(m)})$:

$$m + \frac{4\Gamma(2-d)}{\Gamma(d)} \sum_{k=1}^m (m-k)k^{2d-1}.$$

We then approximate the summation on the right side of this equation by an integral from 1 to m , which leads to the approximation

$$(7) \quad \ddot{V}(s_v^{(m)}) = \frac{2\Gamma(2-d)/\Gamma(d)}{d(2d+1)} m^{2d+1}.$$

Following the same line of reasoning, we also have an approximation for the variance of $h_v^{(m)}$.

$$\ddot{V}(h_v^{(m)}) = \sigma_h^2 \frac{\Gamma(1-d)/\Gamma(d)}{d(2d+1)} m^{2d+1}.$$

The result of Equation 7 is consistent with Theorem 2.2 of Beran (1994), derived using the power spectrum. The theorem states that if the autocorrelation function of a long-range dependent series with fractional difference power d , converges to ck^{2d-1} as k gets large, then the m -blocksum variance converges to

$$\frac{c}{d(2d+1)} m^{2d+1}.$$

In our case $c = 2\Gamma(2-d)/\Gamma(d)$ for s_u and $c = \sigma_h^2\Gamma(1-d)/\Gamma(d)$ for h_u . Let

$$\check{s}_v^{(m)} = \frac{s_v^{(m)}}{m^{d+0.5}}, \quad \text{and} \quad \check{h}_v^{(m)} = \frac{h_v^{(m)}}{m^{d+0.5}\sigma_h}.$$

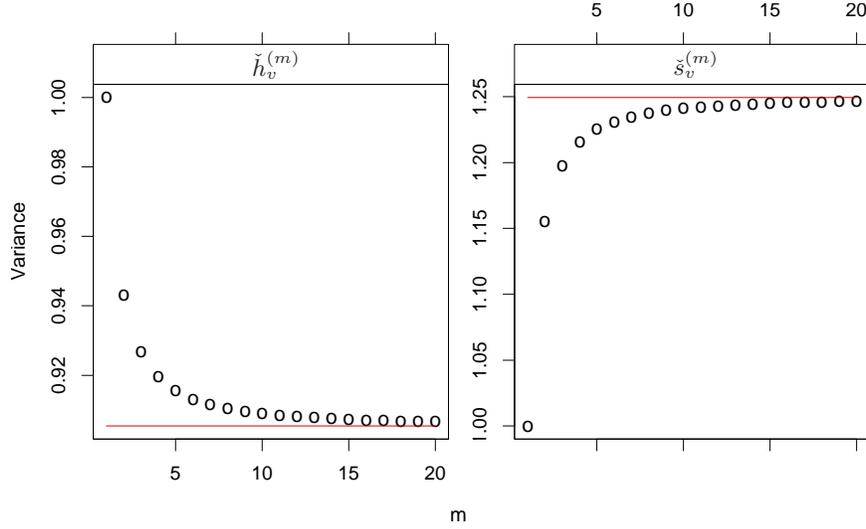


FIG 14. Variances of $\check{h}_v^{(m)}$ and $\check{s}_v^{(m)}$ versus m . Exact variance (\circ) and approximate variance ($—$).

Thus we have

$$\ddot{V}(\check{s}_v^{(m)}) = \frac{2\Gamma(2-d)/\Gamma(d)}{d(2d+1)},$$

$$\ddot{V}(\check{h}_v^{(m)}) = \frac{\Gamma(1-d)/\Gamma(d)}{d(2d+1)}.$$

In Figure 14, the exact variances of $\check{h}_v^{(m)}$ and $\check{s}_v^{(m)}$ are plotted as \circ . The approximate variances, two constant values, are plotted as $—$. The approximation, as we will see in Figure 14, is excellent for $m \geq 3$ for both series. The scaled m -blocksums have nearly constant variances for $m \geq 3$.

8.3.2. *Autocovariance of $s_v^{(m)}$ and $h_v^{(m)}$.* We begin with the autocovariance function of the m -blocksums $s_v^{(m)}$. For $k \geq 1$,

$$C_{s_v^{(m)}}(k) = \sum_{j=-(m-1)}^{m-1} (m - |j|) \rho_s(km + j).$$

Again we use $\check{\rho}_s(k) = \{2\Gamma(2-d)/\Gamma(d)\}k^{2d-1}$ to approximate $\rho_s(k)$, and use integral to approximate summation in the autocovariance function, which leads to

$$\ddot{C}_{s_v^{(m)}}(k) = \frac{2\Gamma(2-d)}{\Gamma(d)} m^{2d+1} k^{2d-1}.$$

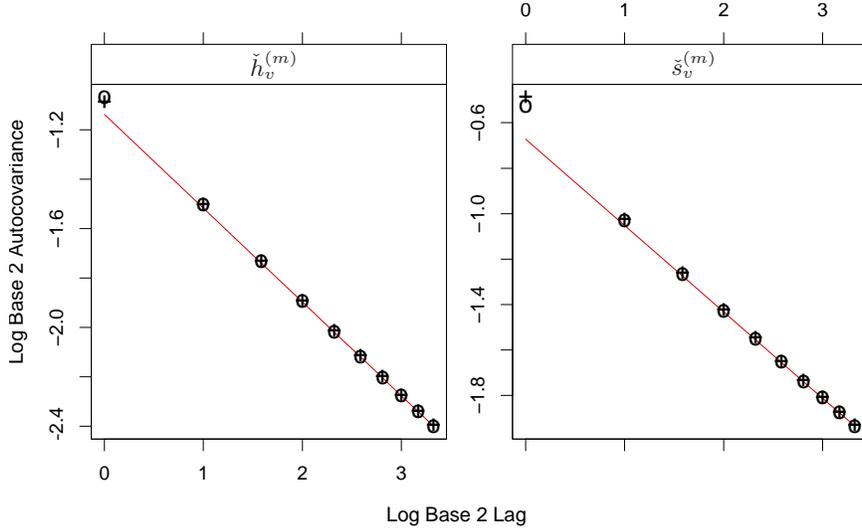


FIG 15. Log of autocovariances of $\check{h}_v^{(m)}$ and $\check{s}_v^{(m)}$ versus $\log_2(k)$ for $m = 2$ (+), $m = 3$ (o), and $\log_2(\check{\rho}_s(k))$ and $\log_2(\check{\rho}_h(k))$ (—),

Follow the same line of reasoning, we have the autocovariance of $h_v^{(m)}$ approximated by

$$\ddot{C}_{h_v^{(m)}}(k) = \sigma_h^2 \frac{\Gamma(1-d)}{\Gamma(d)} m^{2d+1} k^{2d-1}.$$

Because $\rho_s(k) = C_{\check{s}_v^{(1)}}(k)$ is the autocovariance of $\check{s}_v^{(1)}$, and $\rho_h(k) = C_{\check{h}_v^{(1)}}(k)$ is the autocovariance of $\check{h}_v^{(1)}$,

$$\ddot{C}_{\check{s}_v^{(m)}}(k) = \ddot{C}_{\check{s}_v^{(1)}}(k) = \ddot{\rho}_s(k) = \frac{2\Gamma(2-d)}{\Gamma(d)} k^{2d-1}.$$

$$\ddot{C}_{\check{h}_v^{(m)}}(k) = \ddot{C}_{\check{h}_v^{(1)}}(k) = \ddot{\rho}_h(k) = \frac{\Gamma(1-d)}{\Gamma(d)} k^{2d-1}.$$

The scaled m -blocksums $\check{s}_v^{(m)}$ and $\check{h}_v^{(m)}$ have autocovariance functions not dependent on m .

In Figure 15, the exact values of $\log_2(C_{\check{s}_v^{(m)}}(k))$ and $\log_2(C_{\check{h}_v^{(m)}}(k))$ for $m = 2$ (+) and $m = 3$ (o), and $\log_2(\check{\rho}_s(k))$ and $\log_2(\check{\rho}_h(k))$ (—), which is equivalent to $m = 1$, are plotted against $\log_2(k)$. We see the autocovariance functions of $\check{s}_v^{(m)}$ and $\check{h}_v^{(m)}$ are nearly the same for all $m \geq 1$.

For time aggregation with increasing m , only the variances of the scaled m -blocksums $\check{s}_v^{(m)}$ and $\check{h}_v^{(m)}$ change. The autocovariance functions stay nearly constant for $m \geq 1$. Since for $m \geq 3$, the variances of $\check{s}_v^{(m)}$ and $\check{h}_v^{(m)}$ are quite close

to the limit, the processes $\check{s}_v^{(m)}$ and $\check{h}_v^{(m)} \forall m > 3$ are almost the same processes as $\check{s}_v^{(3)}$ and $\check{h}_v^{(3)}$. Both h_u and s_u can be considered as near self-similar process,

9. The Variance-Time Plot: Variances of m Blockmeans $\bar{z}_v^{(m)}$. This section treats the variances of the m -blockmeans $\bar{z}_v^{(m)}$ and how they change with increasing m (time scaling), and with increasing α which causes $\theta(\alpha)$ to increase (source traffic scaling).

As discussed in Section 1, m -blockmean study is a time scaling analysis that has been widely used (e.g., Karagiannis et al. (2004)) as a basis for understanding traffic statistical properties. For example, the variance-time plot, a popular graphical method, is a display of the log m -blockmean variance against log m . Such a plot is linear for a self-similar process. The common pattern in Internet traffic studies (e.g., Fraleigh et al. (2003)) is convex upward with the slope tending to $2H - 2$ where $H = d + 0.5$ is the Hurst parameter. This and other patterns will be seen in this section along with mathematical derivations that describe the patterns. The derivations reveal the properties of the patterns including the result that source traffic scaling, not as widely studied, is a critical factor that interacts with time scaling.

The m -blockmeans of the above three series are related by

$$\bar{z}_v^{(m)} = (1 - \theta(\alpha))\bar{s}_v^{(m)} + \theta(\alpha)\bar{n}_v^{(m)}.$$

Because s_u and n_u are independent, the variances are related by

$$(8) \quad V(\bar{z}_v^{(m)}) = (1 - \theta(\alpha))V(\bar{s}_v^{(m)}) + \theta(\alpha)V(\bar{n}_v^{(m)}).$$

Since n_u is Gaussian white noise with variance 1, $V(\bar{n}_v^{(m)}) = m^{-1}$. So it is quite clear from the beginning that scaling properties are determined by the changing relative contributions of $\sqrt{1 - \theta(\alpha)}s_u$ and $\sqrt{\theta(\alpha)}n_u$ to the variability of the m -blockmeans as m changes and as $\theta(\alpha)$ changes with the traffic rate α .

9.1. Variance of m -Blockmeans.

9.1.1. *Approximation.* In Section 8 we have the approximation $\ddot{V}(s_v^{(m)})$ from Equation 7, which results in an approximation of $V(\bar{z}_v^{(m)})$:

$$(9) \quad \ddot{V}(\bar{z}_v^{(m)}) = (1 - \theta(\alpha))\frac{2\Gamma(2 - d)/\Gamma(d)}{d(2d + 1)}m^{2d-1} + \theta(\alpha)m^{-1}.$$

9.1.2. *Validation.* Validation study for the above variance derivations is an assessment of the accuracy of the approximation of $V(\bar{z}_v^{(m)})$ by $\ddot{V}(\bar{z}_v^{(m)})$, and an assessment of how well $\ddot{V}(\bar{z}_v^{(m)})$ fits sample-variance estimates $\hat{V}(\bar{z}_v^{(m)})$ of $V(\bar{z}_v^{(m)})$

for each of the packet trace segments. Variance time plots are used in the assessments.

For the approximation study, Figure 16 shows results for 3 values of $\theta(\alpha)$: 0.6, 0.8, and 0.975; the estimates of $\theta(\alpha)$ for the the 144 Auckland traces range from about 0.6 to about 0.975. d was taken to be 0.31; it is the median of the estimates from the trace segments and variation of the estimates from this value is small. Let \log_2 be log base 2. In Figure 16, $\log_2\{V(\bar{z}_v^{(m)})\}$ (\circ) and $\log_2\{\ddot{V}(\bar{z}_v^{(m)})\}$ ($—$) are plotted against 15 values of $\log_2(m)$ from 0 to 14. $V(\bar{z}_v^{(m)})$ is computed from Equations 6 and 8. $\ddot{V}(\bar{z}_v^{(m)})$ is computed from Equation 9. The oblique lines will be described shortly. Figure 16 shows that the approximation is excellent.

9.1.3. *Properties of the Variance-Time Plot.* The formula in Equation 9 for the approximate m -blockmean variance, $\ddot{V}(\bar{z}_v^{(m)})$, provides much insight into the

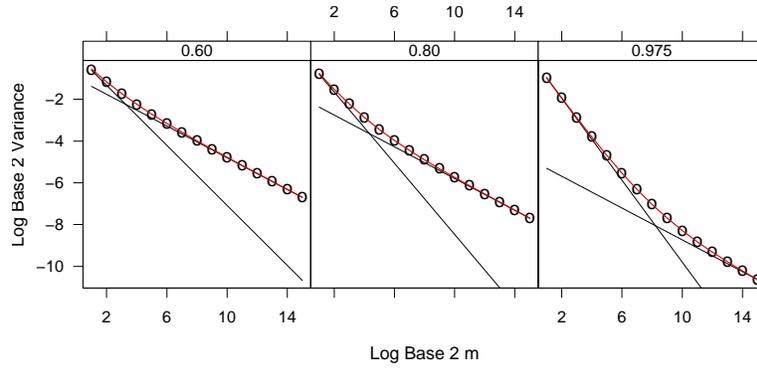


FIG 16. Variance-time plot for 3 values of $\theta(\alpha)$, shown in the strip labels at the top of each panel. The exact $\log_2 V(\bar{z}_v^{(m)})$ (\circ) and the approximate $\log_2 \ddot{V}(\bar{z}_v^{(m)})$ ($—$) are plotted against against $\log_2 m$.

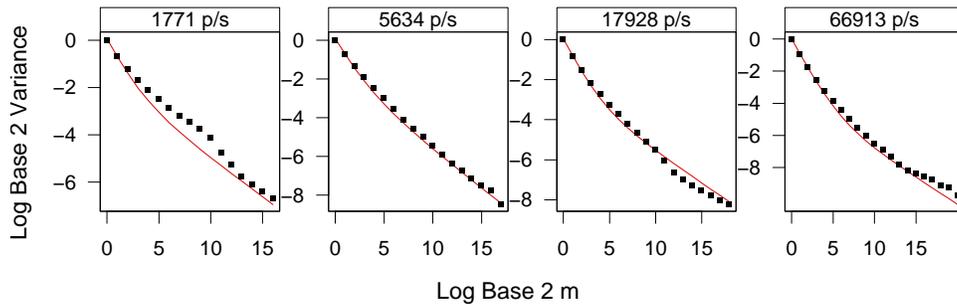


FIG 17. Variance-time plots for four trace segments. The approximate $\log_2 \ddot{V}(\bar{z}_v^{(m)})$ ($—$) and the sample estimates $\log_2 \hat{V}(\bar{z}_v^{(m)})$ (\bullet) against $\log_2 m$.

properties of $\log_2\{\ddot{V}(\bar{z}_v^{(m)})\}$, and therefore the properties of $\log_2\{V(\bar{z}_v^{(m)})\}$, as a function of $\ell = \log_2(m)$. Quite important is the derivative of the functional dependence, the slope on the variance-time plot. If $\theta(\alpha)$ were 1, then $z_u = n_u$, the white noise component, the slope would be everywhere -1 . If $\theta(\alpha)$ were 0, then $z_u = s_u$, the discrete near-fGn component, the slope would be everywhere $(2d - 1) > -1$. While the slope on a variance-time plot for our trace segments, is never everywhere -1 or $2d - 1$, the slope does vary with m and $\theta(\alpha)$ across the traces between these extremes, and under certain circumstances gets close to these values. The slope at $\log_2(m)$ conveys information about the time dependence of the m -blockmean. The closer the slope is to -1 , the closer the m -blockmean is to independence. The closer the slope is to $2d - 1$, the closer the m -blockmean is to a discrete near-fGn process.

We can derive an expression for the slope. Let

$$\phi(d) = \frac{2\Gamma(2-d)}{d(2d+1)\Gamma(d)}.$$

Treating ℓ as a continuous variable and proceeding formally, the slope is

$$\Delta(m, \theta(\alpha), d) = \frac{\partial \log_2 \ddot{V}(\bar{z}_v^{(m)})}{\partial \ell} = \frac{(1 - \theta(\alpha))(2d - 1)\phi(d) - \theta(\alpha)m^{-2d}}{(1 - \theta(\alpha))\phi(d) + \theta(\alpha)m^{-2d}}.$$

For the study of the fit to the trace data, Figure 17 shows results for the 4 Auckland traces described in Section 4 whose traffic rates α span the range of values of observed rates. Analysis, including the method of transforming t_u to the Gaussian image to create observed z_u , was the same as that of Section 4. For each trace, we used the first $n = 2^b$ values of z_u where 2^b is the largest power of 2 less than or equal to the number of observations. The estimate $\hat{V}(\bar{z}_v^{(m)})$ is the sample variance of the $z_v^{(m)}$ for $\log_2(m)$ from 0 to $b - 5$. The values of $\ddot{V}(\bar{z}_v^{(m)})$ are computed using the estimates of d and $\theta(\alpha)$ described in Section 4. In Figure 17, $\log_2\{\hat{V}(\bar{z}_v^{(m)})\}$ (\bullet) and $\log_2\{\ddot{V}(\bar{z}_v^{(m)})\}$ ($-$) are plotted against $\log_2(m)$. The approximate model estimates do an excellent fitting the sample-variance estimates. Since $0 < d < 0.5$, the slope is everywhere negative. As ℓ increases, the slope tends to $2d - 1$. For the value $d = 0.31$ used in our mathematical investigations, $2d - 1 = -0.38$. In each panel of Figure 16, the oblique line through the point for the largest value of m in each panel has slope -0.38 . We can see that in each case, the largest values of m have slopes very close to this value. The second derivative with respect to ℓ is

$$\frac{\partial^2 \log_2 \ddot{V}(\bar{z}_v^{(m)})}{\partial \ell^2} = \frac{\log(2)4d^2\theta(\alpha)(1 - \theta(\alpha))\phi(d)m^{-2d}}{((1 - \theta(\alpha))\phi(d) + \theta(\alpha)m^{-2d})^2} > 0,$$

so the slope increases monotonically with ℓ , and $\log_2\{\ddot{V}(\bar{z}_v^{(m)})\}$, as a function of ℓ is convex, which is the standard empirically-observed pattern.

In addition,

$$\frac{\partial^2 \log_2 \ddot{V}(\bar{z}_v^{(m)})}{\partial \ell \partial \theta(\alpha)} = \frac{-2d\phi(d)m^{-2d}}{\{(1 - \theta(\alpha))\phi(d) + \theta(\alpha)m^{-2d}\}^2} < 0.$$

Thus for fixed d and m , the slope decreases with $\theta(\alpha)$. This can be seen in Figure 16. For example, as $\theta(\alpha)$ increases, there is an increase in the value of m for which the maximum slope of -0.38 is nearly achieved. For the three values of $\theta(\alpha)$ in Figure 16, the slopes at $m = 1$ are $\Delta(1, (0.6, 0.8, 0.975), 0.31) = (-0.72, -0.85, -0.98)$, which decrease with $\theta(\alpha)$; in the three panels of Figure 16, the slopes of the oblique lines through the points for $m = 1$ have these three derived values.

The m -blockmeans across m and $\theta(\alpha)$ vary from nearly independent processes for small m and large $\theta(\alpha)$ to near-fGn process when m gets large for all values of $\theta(\alpha)$. The m at which the m -blockmean becomes near-fGn increases as $\theta(\alpha)$ increases. $\theta(\alpha)$ increases with the traffic rate α . This dependence of the statistical properties of m -blockmean variances on the traffic rate for interarrivals has not been previously recognized.

10. Multifractal Image Autocorrelation Approximations: Validation and Properties. This section describes approximations for autocorrelation of the multifractal image. There are two purposes. One is to provide simpler descriptions that contribute to our fundamental understanding of the drivers of the traffic statistical properties. The fundamentals are described in Section 12 that provides insights about the properties of traffic statistics; we do this even for cases where we have exact formulas. A second purpose is to provide approximate formulas in cases where we are unable to derive exact formulas; the approximate formulas are validated just as we would an exact one.

10.1. *Approximation of the Weibull MFSD by the Multiplicative MFSD.* The multiplicative (log normal) MFSD discussed in Section 1 is simpler than the Weibull MFSD because the transformation to the multifractal image is simpler:

$$t_u = \exp\{\tau(\sqrt{1 - \theta(\alpha)}s_u + \sqrt{\theta(\alpha)}n_u) + \mu\}.$$

Because $t_u t_{u-k}$ is also log normal, $L(2\mu, 2\tau^2(1 + \rho_z(k)))$, the autocorrelation function of the multiplicative MFSD has a simple formula

$$(10) \quad \dot{\rho}_t(k) = \frac{e^{\tau^2 \rho_z(k)} - 1}{e^{\tau^2} - 1}.$$

There is no simple closed-form formula for the autocorrelation, $\rho_t(k)$, of the Weibull MFSD, but we found through simulation that $\dot{\rho}_t(k)$ provides an excellent

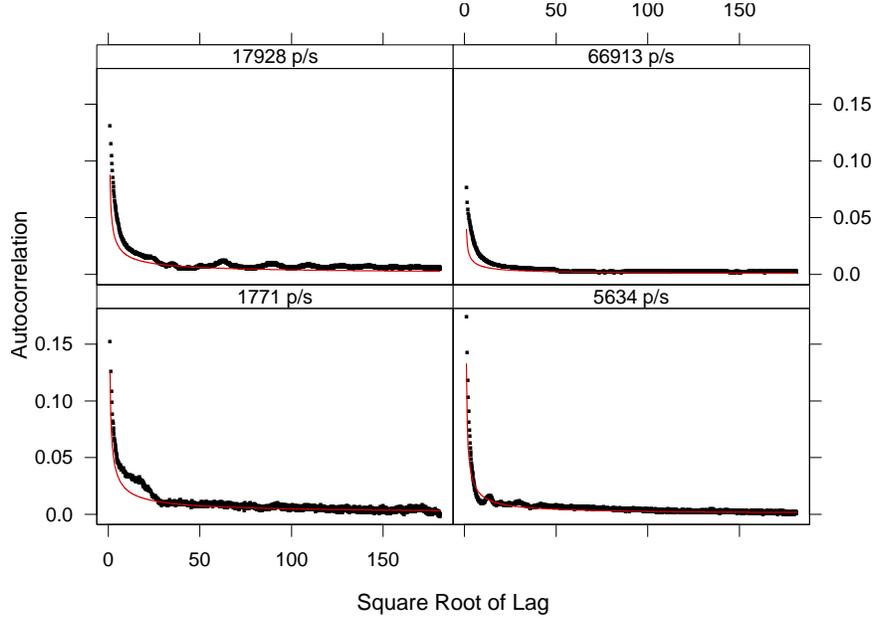


FIG 18. Nonparametric estimates of autocorrelation (\bullet) and estimates from fitted multiplicative MFSD models ($-$).

approximation of $\rho_t(k)$. We also found that the standard nonparametric estimate of $\rho_t(k)$ from each of the trace segments is very well matched by $\hat{\rho}_t(k)$ with $d = 0.31$ and μ and τ^2 estimated from each segment. The estimates $\hat{\mu}$ and $\hat{\tau}^2$ are those for which the first two moments of the log normal match those of the Weibull with parameters equal to $\hat{\lambda}(\alpha)$ and α , those described in Section 3. This is illustrated in Figure 18. The nonparametric estimates (\bullet) and the estimates from $\hat{\rho}_t(k)$ ($-$) are plotted against lag for the 4 trace segments used in previous sections. Taking the statistical variability into account, which includes correlation in the estimates across the lags, the autocorrelation estimates for the multiplicative MFSD is a very good approximation of the of the nonparametric estimates.

10.2. *A Power Series Approximation.* The autocorrelations of z_u and t_u , while highly persistent, are not large; almost all values are below 0.25, and beyond the first few lags, are below 0.10. Furthermore, we found that estimates of τ^2 across all traces range from 0.7 to 1.4. The resulting values of $\tau^2 \rho_z(k)$ are small enough that the term $\exp\{\tau^2 \rho_z(k)\}$ in Equation 10 is very well approximated by a first order power series approximation, resulting in the even simpler autocorrelation approximation

$$\hat{\rho}_t(k) = \frac{\tau^2 \rho_z(k)}{e^{\tau^2} - 1}.$$

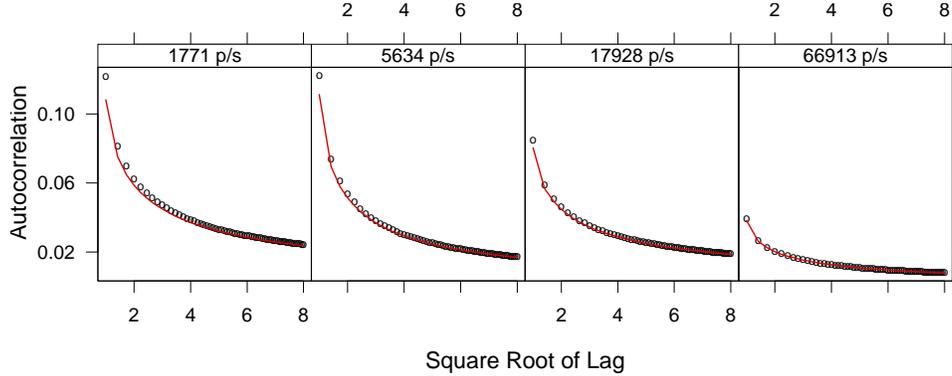


FIG 19. Power series approximation (—) of multiplicative MFSD autocorrelations (\circ).

Since $\rho_z(k)$ is always positive we know that $\ddot{\rho}_t(k)$ will be somewhat smaller than $\dot{\rho}_t(k)$. Figure 19 graphs the values of $\dot{\rho}_t(k)$ from Figure 18 for lags 1 to 64 (\circ) and the values of $\ddot{\rho}_t(k)$ (—). Except for lag 1, which has a minor departure, the approximation is very close. The astonishing result, discussed next, is that the autocorrelations of the multifractal image are very close to being proportional to those of the Gaussian image.

10.3. *Heuristics for the Multifractal Image.* The mechanism for going to the multifractal image, the nonlinear transformation, is simple to describe. However, the general result is a dramatic change in the statistical properties, from Gaussian to highly nonlinear. The statistics of t_u , without the benefit of the simple transformation mechanism of the model, have an appearance of a highly complex mechanism.

However, the very surprising result is that the form of the second moments are, to a very good approximation, preserved up to a multiplicative constant under the transformation. This means that the properties of the autocorrelations of the Gaussian image, studied extensively in previous and coming sections, hold to a good approximation for the multifractal image. This includes the time-aggregation and source-aggregation properties studied next in Sections 9 and 6.

11. A Derived Model for Time Aggregates of t_u . It is common in mathematical studies of engineering performance and control to use time aggregates of packet arrivals as source traffic inputs and to assume a long-range dependent Gaussian process; this can result in a more tractable mathematical construct because of the Gaussian assumption. Time aggregates are also used in simulations to speed computation (e.g., Baiocchi and Vacirca (2007); Bäuerle and Rieder (2000); Carofiglia et al. (2007); Gu et al. (2004); Kim and Shroff (2001); Liu et al. (2004);

[Kiddle et al. \(2003\)](#)). However, there is no guarantee that time aggregates of a non-linear process tend to Gaussian.

In this section we verify that aggregates — the m -blockmeans, $\bar{t}_v^{(m)}$, of t_u — do tend to Gaussian, and we derive and validate an variance and an autocovariance function for the $\bar{t}_v^{(m)}$. This not only validates the Gaussian assumption but provides a more incisive description of the autocovariance function that can provide stronger results than a general long range dependence assumption frequently used in the past.

11.1. *Autocorrelation Function of $\bar{z}_v^{(m)}$* . First we examine the autocorrelation function $\rho_{\bar{z}_v^{(m)}}(k)$ of the m -blockmeans of z_u , and how it changes with increasing m or increasing $\theta(\alpha)$. Previous studies have investigated these properties empirically. [Feldmann et al. \(1998b\)](#); [Gilbert et al. \(1999\)](#) observed that the energy of wavelet models for counts increases as the time aggregation increases. [Hannig et al. \(2001\)](#) observed that the autocorrelation of arrival counts in fixed intervals increases with increasing time aggregation. The heuristics that guide our understanding of the properties of the m -blockmean variances enable us to readily predict these outcomes as m increases, and to add predictions of outcomes as the traffic rate α increases. As m increases for fixed $\theta(\alpha)$, then the low-pass filtering of the m -blockmean acting on z_u removes a larger fraction of the variance of n_u than of s_u , and $\rho_{\bar{z}_v^{(m)}}(k)$ increases. For fixed m , as α increases, $\theta(\alpha)$ increases; this increases the contribution of the white noise component, $\theta(\alpha)\bar{n}_v^{(m)}$, to the variance of $\bar{z}_v^{(m)}$, reducing $\rho_{\bar{z}_v^{(m)}}(k)$.

We can calculate the $\rho_{\bar{z}_v^{(m)}}(k)$ from the GFSD model and check these properties. We have

$$\rho_{\bar{z}_v^{(m)}}(k) = \frac{C_{\bar{z}_v^{(m)}}(k)}{V(\bar{z}_v^{(m)})} = \frac{(1 - \theta(\alpha))C_{\bar{s}_v^{(m)}}(k)}{V(\bar{z}_v^{(m)})},$$

where $V(\bar{z}_v^{(m)})$ is calculated from Equations 6 and 8, and

$$C_{\bar{s}_v^{(m)}}(k) = m^{-1} \sum_{j=-(m-1)}^{m-1} \left(1 - \frac{|j|}{m}\right) \rho_s(km + j).$$

Figure 20 graphs $\log_2\{\rho_{\bar{z}_v^{(m)}}(k)\}$ against $\log_2(k)$ for $d = 0.31$, $\theta(\alpha) = (0.60, 0.80, 0.975)$, and $m = 1$ (—), $m = 50$ (—), and 3000 (—). Figure 21 graphs $\log_2\{\rho_{\bar{z}_v^{(m)}}(k)\}$ against $\log_2(k)$ for $d = 0.31$, $m = (1, 50, 3000)$, and $\theta(\alpha) = 0.975$ (—), $\theta(\alpha) = 0.80$ (—), and 0.60 (—). The figures show that the autocorrelation increases with increasing m or with decreasing $\theta(\alpha)$.

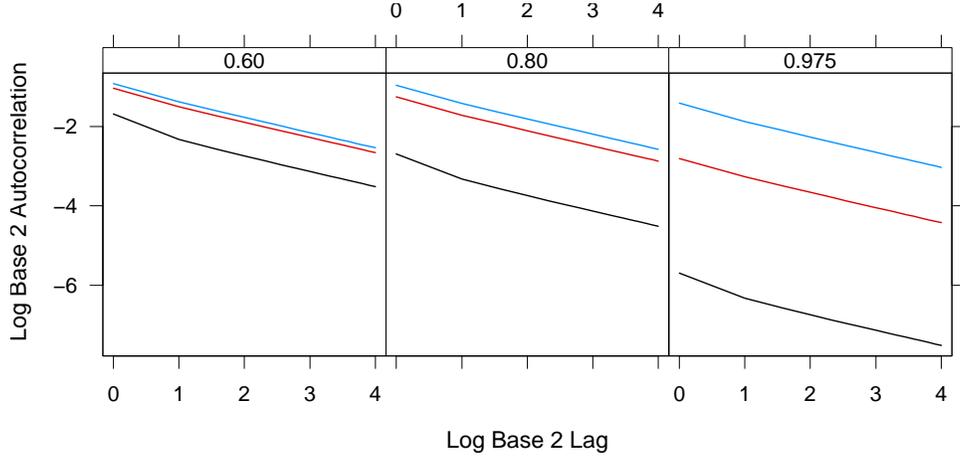


FIG 20. $\log_2\{\rho_{z_v^{(m)}}(k)\}$ is plotted against $\log_2(k)$ for $m = 1$ (—), $m = 50$ (—), $m = 3000$ (—) for 3 values of $\theta(\alpha)$.

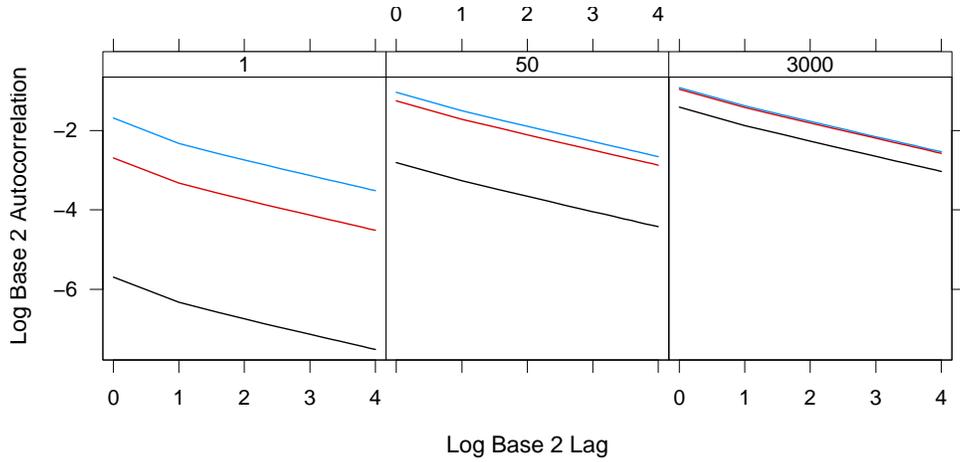


FIG 21. $\log_2\{\rho_{z_v^{(m)}}(k)\}$ is plotted against $\log_2(k)$ for $\theta(\alpha) = 0.975$ (—), $\theta(\alpha) = 0.80$ (—), $\theta(\alpha) = 0.60$ (—) for 3 values of m .

11.2. *Limiting Behavior of Partial Sums of t_u .* In the MFSD model, the inter-arrivals t_u are a monotone function of the GFSD z_u . Past work has a number of informative limiting results of time aggregates of functionals of Gaussian long-range dependent processes (Dobrushin and Major, 1979; Doukhan et al., 2003; Surgailis, 2000; Taqqu, 1975, 1977). We will apply the results of Taqqu (1975) to the multiplicative MFSD, which we think of as a good approximation of the true Weibull MFSD. Next, we describe properties of the GFSD and the multiplicative MFSD

that allow us to invoke results from this reference that a partial-sum of the multiplicative MFSD goes to fractional Brownian motion (fBm) whose stationary increment is fGn.

The GFSD z_u is a Gaussian long-range dependent process with mean 0 and variance 1. In addition, in Section ??, we have shown that the autocorrelation of z_u approaches $2(1 - \theta)\Gamma(2 - d)/\Gamma(d)k^{2d-1}$ as $k \rightarrow \infty$. The multiplicative MFSD with mean centered at 0 is

$$G(z_u) = e^{\tau z_u + \mu} - e^{\mu + \tau^2/2},$$

where μ and τ^2 are the mean and variance of the log multiplicative MFSD. $G(z_u)$ has variance $(e^{\tau^2} - 1)e^{2\mu + \tau^2}$.

The Hermite polynomial of order k is

$$H_k(z) = (-1)^k e^{z^2/2} \frac{d^k}{dz^k} (e^{-z^2/2}).$$

The first three are $H_0(z) = 1$, $H_1(z) = z$, and $H_2(z) = z^2 - 1$. The Hermite expansion of $G(z_u)$ is

$$G(z_u) = \sum_{k=0}^{\infty} \frac{b_k}{k!} H_k(z_u),$$

where

$$b_k = E\{H_k(z_u)G(z_u)\}.$$

The Hermite rank is defined as the smallest k in the expansion for which $b_k \neq 0$. For $G(z_u)$, $b_0 = 0$ and

$$b_1 = E\{z_u(e^{\tau z_u + \mu} - e^{\mu + \tau^2/2})\} = \tau e^{\tau^2/2 + \mu}.$$

So the Hermite rank of $G(z_u)$ is 1. Let

$$\delta_m = \sqrt{(\tau^2 e^{\tau^2 + 2\mu}) \frac{2(1 - \theta)\Gamma(2 - d)/\Gamma(d)}{d(1 + 2d)} m^{1+2d}}.$$

Using the above facts about z_u and $G(z_u)$, we can now invoke the results of [Taqqu \(1975\)](#). Let $0 \leq w \leq 1$, and let $[m\omega]$ be the greatest integer in $m\omega$. The scaled partial sums

$$\frac{1}{\delta_m} \sum_{u=1}^{[m\omega]} G(z_u)$$

converge weakly to fBm as $m \rightarrow \infty$.

This result is not a proof that the $\bar{t}_v^{(m)}$ as process in m tends to Gaussian but it does make it highly plausible. However, extensive empirical study of the $\bar{t}_v^{(m)}$

provides a convincing validation that the $\bar{t}_v^{(m)}$ are Gaussian for small values of m . For simulated i.i.d Weibull random numbers with $\lambda = 0.7$, the lower end of the parameter estimates, $\bar{t}_v^{(m)}$ approaches Gaussian with $m \geq 70$. For simulated i.i.d Weibull random numbers with $\lambda = 1$, the higher end of the parameter estimates, $\bar{t}_v^{(m)}$ approaches Gaussian with less aggregation, $m \geq 50$. For the lowest rate live traces, it needs $m \geq 20$ for $\bar{t}_v^{(m)}$ to approach Gaussian, which is equivalent to aggregation over 10ms intervals on average. m needs to be larger for high rate traces. For the highest rate traces, it needs $m \geq 120$ for $\bar{t}_v^{(m)}$ to approach Gaussian, which is equivalent to aggregation over 1ms–2ms intervals on average.

11.3. *Variance and Autocovariance of $\bar{t}_v^{(m)}$* . For sufficient time aggregation, we can use a Gaussian process as a model for the $\bar{t}_v^{(m)}$. In this Section we provide the variance and covariance function of $\bar{t}_v^{(m)}$. Hence we offer a concrete model for the time aggregates.

Let $\sigma_t^2 = V(t_u)$. Variance of $\bar{t}_v^{(m)}$ equals to

$$V(\bar{t}_v^{(m)}) = \sigma_t^2 \left\{ m^{-1} + 2m^{-1} \sum_{k=1}^m \left(1 - \frac{k}{m}\right) \rho_t(k) \right\}.$$

Applying the approximation of t_u autocorrelation in Section 10.2

$$\ddot{\rho}_t(k) = \frac{\tau^2}{e^{\tau^2} - 1} \rho_z(k) = \frac{\tau^2(1 - \theta)}{e^{\tau^2} - 1} \rho_s(k),$$

and $\ddot{\rho}_s(k) = \{2\Gamma(2 - d)/\Gamma(d)\}k^{2d-1}$, we use this to approximate $V(\bar{t}_v^{(m)})$:

$$\sigma_t^2 \left\{ m^{-1} + \frac{\tau^2(1 - \theta)}{e^{\tau^2} - 1} \frac{2\Gamma(2 - d)}{\Gamma(d)} (2m^{-1} \sum_{k=1}^m \left(1 - \frac{k}{m}\right) k^{2d-1}) \right\}.$$

We use an integral from 1 to m to approximate the summation in the above expression, and obtain the following. The parameters are functions of the arrival rate α . The exact expressions of $\sigma_t^2(\alpha)$ and $\tau^2(\alpha)$ are shown later.

$$(11) \quad \ddot{V}(\bar{t}_v^{(m)}) = \sigma_t^2(\alpha) \left\{ \frac{\tau^2(\alpha)(1 - \theta(\alpha))}{e^{\tau^2(\alpha)} - 1} \frac{2\Gamma(2 - d)/\Gamma(d)}{d(2d + 1)} m^{2d-1} + m^{-1} \right\}.$$

Autocovariance of $\bar{t}_v^{(m)}$ equals to

$$C_{\bar{t}_v^{(m)}}(k) = \sigma_t^2 m^{-1} \left\{ \sum_{j=-(m-1)}^{m-1} \left(1 - \frac{|j|}{m}\right) \rho_t(km + j) \right\}.$$

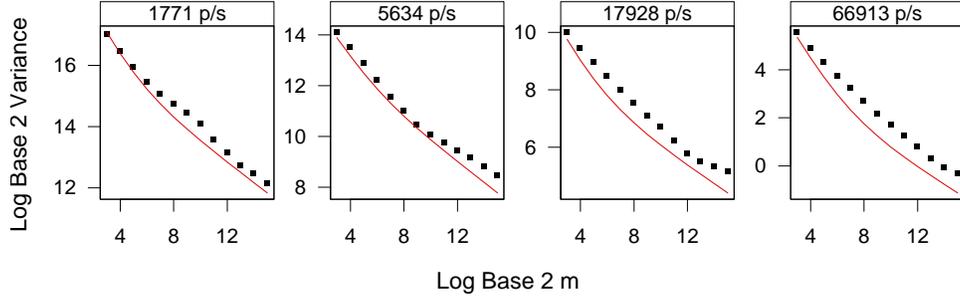


FIG 22. Variance-time plots of $\bar{t}_v^{(m)}$ for four trace segments. The sample estimates $\log_2 \hat{V}(\bar{t}_v^{(m)})$ (●) and the approximate $\log_2 \ddot{V}(\bar{t}_v^{(m)})$ (—) are plotted against $\log_2 m$.

Again applying the approximations $\ddot{\rho}_t(k)$ and $\ddot{\rho}_s(k)$ and using integral to approximate summation, we have

$$(12) \quad \ddot{C}_{\bar{t}_v^{(m)}}(k) = \sigma_t^2(\alpha) \frac{\tau^2(\alpha)(1 - \theta(\alpha))}{e^{\tau^2(\alpha)} - 1} \frac{2\Gamma(2-d)}{\Gamma(d)} m^{2d-1} k^{2d-1}.$$

The parameters in Equations 11 and 12 depend only on the rate α . Through matching the first two moments of log normal with the first two moments of Weibull we have

$$\tau^2(\alpha) = \log \Gamma \left(1 + \frac{2}{\lambda(\alpha)} \right) - 2 \log \Gamma \left(1 + \frac{1}{\lambda(\alpha)} \right).$$

The variance $\sigma_t^2(\alpha)$ of Weibull is

$$\sigma_t^2(\alpha) = \frac{\Gamma(1 + \frac{2}{\lambda(\alpha)})}{\alpha^2 \Gamma^2(1 + \frac{1}{\lambda(\alpha)})} - \frac{1}{\alpha^2}.$$

$\lambda(\alpha)$ and $\theta(\alpha)$ follow the logistic models.

A Gaussian process with the approximate variance and autocovariance function specified by Equations 11 and 12 is a concrete model for the $\bar{t}_v^{(m)}$, whose coefficients depend on the rate α only, as the MFSD and GFSD models. Furthermore the ratio of Equations 11 and 12 is an approximation of the $\bar{t}_v^{(m)}$ autocorrelation, which behaves as fGn autocorrelation $d(2d+1)k^{2d-1}$ as $m \rightarrow \infty$.

In Figure 22, the sample estimates $\log_2 \{\hat{V}(\bar{t}_v^{(m)})\}$ (●) and the model approximate $\log_2 \{\ddot{V}(\bar{t}_v^{(m)})\}$ (—) are plotted against $\log_2(m)$ for the four selected traces. This is the variance time plot for the multifractal image. The approximate model estimates do a good fitting the sample-variance estimates. The value of m is from 2^3 to 2^{15} . For the trace with rate 1771 p/s, it is equivalent to aggregation from 4ms to 18sec. For the trace with rate 66913 p/s, it is equivalent to aggregation from

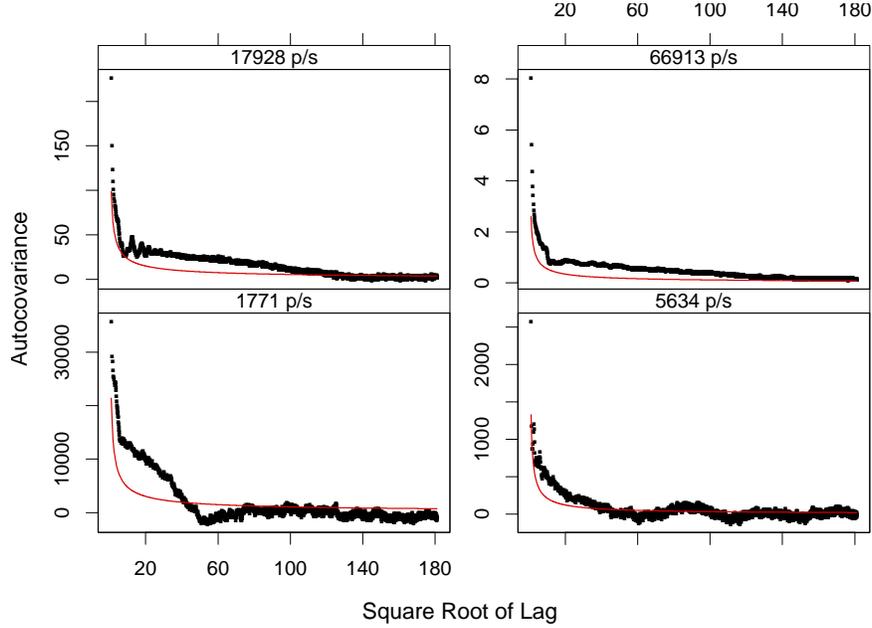


FIG 23. The sample estimates of autocovariance $\hat{C}_{\hat{t}_v^{(m)}}(k)$ (\bullet) and the approximate autocovariance $\check{C}_{\hat{t}_v^{(m)}}(k)$ (---) are plotted against \sqrt{k} for $m = 25$, for four trace segments.

0.1ms to 0.5sec. The same intuition obtained from the Gaussian image variance time plot applies to the multifractal image as well.

In Figure 23 the sample estimates $\hat{C}_{\hat{t}_v^{(m)}}(k)$ (\bullet) and the approximate autocovariance $\check{C}_{\hat{t}_v^{(m)}}(k)$ (---) are plotted against the square root of lag k . The approximate autocovariance specified by Equation 12 provides an excellent fitting to the sample estimates for all the live traces for $m \geq 25$, a small number. Figure 23 shows the four selected traces. The variance of interarrivals, $\sigma_t^2(\alpha)$, approaches 0 quickly as α increases. Hence the variance and autocovariance of $\hat{t}_v^{(m)}$ for fixed m decrease fast with α .

12. Discussion: Fundamentals of Traffic Statistical Properties. The basic characteristic from which mathematical tractability and strong intuition arise from our modeling of Internet traffic is the form of the GFSD model for the Gaussian image z_u . The model is an additive mixture of two components: a near-fGn component $\sqrt{1 - \theta(\alpha)}s_u$ with variance $1 - \theta(\alpha)$, and a white noise component $\sqrt{\theta(\alpha)}n_u$ with variance $\theta(\alpha)$. The term “near-fGn” is used because the autocorrelation function of s_u is very well approximated by $\Gamma(1 - d)/\Gamma(d)k^{2d-1}$ for lag $k = 3$ and higher.

Fundamental understanding of the statistical properties of the traffic is based on (1) simple formulas arising from the simple form of the GFSD model; and (2) the effect of $\sqrt{1 - \theta(\alpha)}s_u$ and $\sqrt{\theta(\alpha)}n_u$ on traffic statistics arising from operations on z_u . In this article we (1) present formulas, some exact and some approximate; and (2) investigate a number of operations: transformation to the multifractal image t_u , m -blocksums, increases in the packet arrival rate α , and partial sums of t_u . In general, mathematical investigations of traffic statistics using the MFSD for t_u can proceed in this way.

12.1. *Transformation from z_u to t_u .* The transformation of z_u to the multifractal image t_u results in a Weibull marginal distribution for t_u . However, to a good approximation, the process is a multiplicative MFSD

$$t_u \approx \exp\{\mu + \tau(\sqrt{1 - \theta(\alpha)}s_u + \sqrt{\theta(\alpha)}n_u)\}.$$

The two components of the GFSD act multiplicatively in t_u . This implies the marginal distribution of t_u is log normal, but this not far off from the Weibull for the range of estimates of the shape λ from the trace segments.

Even more, the autocorrelation function arising from this multiplicative MFSD is to a very good approximation a constant times the autocorrelation function of z_u . This fundamental characteristic is a fortuitous interpretive bonus because, with adjustments for the constant, statistical properties derived for z_u that involve second moments can be carried over to t_u . The good fortune occurs because the values of τ and $1 - \theta(\alpha)$ that arise in practice are sufficiently small that an exponential in the autocorrelation formula for the multiplicative MFSD can be well approximated by a first order power series.

12.2. *Increase in the Traffic Rate α .* As α increases, $\lambda(\alpha)$ and $\theta(\alpha)$ go to 1. This is demonstrated by mathematical derivations based on the MFSD, and estimates based on the trace segments. Simple logit models account for the change. Let $\theta^*(\alpha) = 2^{-7.209}\alpha^{0.746}$ and $\lambda^*(\alpha) = 2^{-5.356}\alpha^{0.628}$, then the models are

$$\begin{aligned}\lambda(\alpha) &= \lambda^*(\alpha)/(1 + \lambda^*(\alpha)) \\ \theta(\alpha) &= \theta^*(\alpha)/(1 + \theta^*(\alpha)).\end{aligned}$$

The limits mean that the marginal distribution of t_u tends to exponential; the variance, $1 - \theta(\alpha)$, of the near-fGn component tends to 0; the variance, $\theta(\alpha)$, of the white noise component tends to 1; and t_u tends to a Poisson process. The autocorrelation functions at lags $k > 0$ for t_u and z_u go to 0 uniformly. Each has the form $(1 - \theta(\alpha))a_k$, where $a_k > 0$ do not depend on $\theta(\alpha)$. From the above logit model for $\theta(\alpha)$, $1 - \theta(\alpha) = 1/(1 + \theta^*(\alpha))$, so the autocorrelations go to 0 like $\alpha^{-0.628}$.

12.3. *m* Blockmeans. Fundamental insight about the effect of $\sqrt{1 - \theta(\alpha)}s_u$ and $\sqrt{\theta(\alpha)}n_u$ the *m*-blockmeans of z_u is straightforward. Because this involves second moments, results hold for t_u as well.

The *m*th block mean of z_u is

$$\bar{z}_v^{(m)} = (1 - \theta(\alpha))\bar{s}_v^{(m)} + \theta(\alpha)\bar{n}_v^{(m)}.$$

The blockmean operation, before taking every *m*th value, is a linear, low-pass, digital filter with transfer function $m^{-1} \sin^2(\pi m f) \sin^{-2}(\pi f)$. The filter removes almost all frequencies outside of an interval $[0, f^*(m)]$ cycles/interarrival in which $f^*(m) \rightarrow 0$ as *m* increases. $\sqrt{1 - \theta(\alpha)}s_u$ has much of its total power near 0 frequency, whereas the power for $\sqrt{\theta(\alpha)}n_u$ is spread uniformly across all frequencies. This means the reduction in variance of $\sqrt{\theta(\alpha)}n_v^{(m)}$ starting at $m = 1$ advances more quickly than that of $(1 - \theta(\alpha))\bar{s}_v^{(m)}$, so $\sqrt{\theta(\alpha)}n_v^{(m)}$ dominates the reduction in the variance of $\bar{z}_v^{(m)}$, which drops to order $(\theta(\alpha)m^{-1})$. Once *m* is sufficiently large, $\sqrt{\theta(\alpha)}n_v^{(m)}$ has little power left relative to $\sqrt{(1 - \theta(\alpha))\bar{s}_v^{(m)}}$, so the latter dominates the reduction in variance of $\bar{z}_v^{(m)}$, which drops to order $(1 - \theta(\alpha))m^{2d-1}$.

However, we can also see that the packet rate α plays a roll because $\theta(\alpha) \rightarrow 1$ as α increases. When $\theta(\alpha)$ is close to 1, for small values of *m*, the variance of $\bar{z}_v^{(m)}$ is mostly due to $\theta(\alpha)\bar{n}_v^{(m)}$ to initially the reduction in variance is like m^{-1} . But for the smallest values of $\theta(\alpha)$ observed in our traces, about 0.6, the drop in variance is affected for small *m* by $\bar{s}_k^{(m)}$ and the decrease is less rapid as *m* increases from 1.

Another immediate insight is that an increasing packet rate α changes the results of the operation. With the increase, $\theta(\alpha)$ increases. This means the variance of the white noise component is larger at the outside, for $m = 1$, so there an increase in the value of *m* for which $\bar{s}_v^{(m)}$ dominates.

This intuition is backed up the quantitative results of the *m*-blockmean operation in Section 9. Consider as an example the variance-time plot. For all values of $\theta(\alpha)$ the slope on the plot tends to $2d - 1$ as *m* increases, that for $\bar{s}_v^{(m)}$. So $\bar{s}_v^{(m)}$ eventually dominates. For the slope at $m = 1$ and for $\theta(\alpha)$ at its smallest estimated value, about 0.6, the slope on the variance time plot for the at $m = 1$ is approximately $-2/3$. At $\theta(\alpha)$ increases, the slope at $m = 1$ tends to -1 , that for an independent process. So for small *m*, $\theta(\alpha)\bar{n}_v^{(m)}$ dominates at $m = 1$ for large $\theta(\alpha)$.

However, this does not mean that it is prudent to assume a Poisson process for mathematical or engineering study large α . The caution arises because the near-fBn component always has a positive variance, and a stochastic process converging to a limit does not mean all operations (functions) of the process converge to the operation on the limit. The next discussion provides and example.

However, the very surprising result is that the form of the second moments are, to a very good approximation, preserved up to a multiplicative constant under the transformation. This means that the properties of the autocorrelations of the Gaussian image, studied extensively in previous and coming sections, hold to a good approximation for the multifractal image. This includes the time-aggregation and source-aggregation properties studied next in Sections 9 and 6.

13. Appendix: Derivations of λ and θ for Changing α .

13.1. *Heuristic Derivation of λ .* We present a heuristic derivation for λ in this section, under the simplified assumption of renewal processes. Assume there are r i.i.d renewal processes, where r is a positive integer. Each renewal process has a Weibull marginal distribution with parameters α and λ . From the analysis of the trace segments in the previous sections, we notice the marginal distribution for the live traffic superposition process is Weibull with parameters $\lambda(r)$ and $\alpha(r)$, where the increasing traffic rate is

$$\alpha(r) = r\alpha.$$

Based on the Weibull marginal distribution for the individual renewal processes, we have the marginal density of their superposition process (Cox, 1962):

$$g^r(t) = -\frac{d}{dt} \left(e^{-(t\alpha\Gamma(1+\frac{1}{\lambda}))^\lambda} \left(\int_t^\infty \alpha e^{-(x\alpha\Gamma(1+\frac{1}{\lambda}))^\lambda} dx \right)^{r-1} \right).$$

The median β^r of the distribution with density $g^r(t)$ is found by solving the following equation:

$$(13) \quad e^{-(\beta^r \alpha \Gamma(1+\frac{1}{\lambda}))^\lambda} \left(1 - \alpha \int_0^{\beta^r} e^{-(x\alpha\Gamma(1+\frac{1}{\lambda}))^\lambda} dx \right)^{r-1} = 0.5.$$

We approximate the distribution with density $g^r(t)$ by a Weibull distribution with parameters $\lambda(r)$ and $\alpha(r)$ whose median matches β^r .

$$(14) \quad \beta^r = \frac{(\log 2)^{\frac{1}{\lambda(r)}}}{r\alpha\Gamma(1 + 1/\lambda(r))}.$$

We solve Equation 14 for $\lambda(r)$. The right hand side of Equation 14 changes monotonically with $\lambda(r)$. Hence there is a unique solution for $\lambda(r)$.

13.2. *Heuristic Derivation of θ .* Assume there are r i.i.d MFSD source processes t_u . Each has a Weibull marginal distribution with parameters α and λ . The corresponding Gaussian image z_u of a source process t_u follows a GFSD model

with parameters θ and d . From the analysis of trace segments in the previous sections, we observe the values of the fractional difference parameter d do not change appreciably under different traffic rates. We then fix the value of d to be the median of the estimates in Section 4. Let

$$d(r) = d = 0.31.$$

For the Weibull marginal distribution of the superposition process t_u^r , we have the estimates $\hat{\alpha}(r) = r\alpha$, and $\hat{\lambda}(r)$ obtained from Equation 14.

$\theta(r)$ can be obtained from the autocorrelation at lag 1 for the Gaussian image z_u^r of the superposition process t_u^r as follows:

$$(15) \quad \theta(r) = 1 - \frac{\rho_z^r(1)(2-d)}{d+1}.$$

To estimate $\theta(r)$, we first compute $\rho_t^r(1)$ for the superposition process t_u^r , then obtain $\rho_z^r(1)$ from $\rho_t^r(1)$, and apply Equation 15.

In order to obtain $\rho_t^r(1)$, first we examine the sources of the arrivals that lead to the two consecutive interarrival times in the superposition process, t_u^r and t_{u+1}^r . Assume the u -th arrival a_u^r in the superposition process is the j th arrival from source 1. Note $t_{u+1}^r = a_{u+1}^r - a_u^r$. Let $t_{i,j}$ be the j th interarrival time from source i , and V_i be a forward recurrence time (the time from an arbitrary time point until the next arrival) for source i . To find the autocorrelation between t_u^r and t_{u+1}^r , we examine the following five cases:

1. The arrivals a_{u-1}^r and a_{u+1}^r both come from source 1. Then t_u^r and t_{u+1}^r are two consecutive interarrival times from source 1. We have $\rho_t^r(1) = \rho_t(1)$.
2. The arrival a_{u+1}^r comes from source 1, but the arrival a_{u-1}^r comes from a different source i . We have $t_{u+1}^r = t_{1,j+1}$. And $t_u^r = V_1 = t_{1,j} - \eta$, where η is the sum of interarrival times and forward recurrence times from sources other than source 1. Then as in Case 1, $\rho_t^r(1) = \rho_t(1)$.
3. The arrival a_{u-1}^r comes from source 1, but the arrival a_{u+1}^r comes from a different source i . Then $t_u^r = t_{1,j}$ and $t_{u+1}^r = V_i$. Since the individual source processes are independent, $\rho_t^r(1) = 0$.
4. The arrivals a_{u-1}^r and a_{u+1}^r come from different sources, and neither comes from source 1. Assume arrival a_{u+1}^r come from source i , $i \neq 1$. Then $t_u^r = V_1$ and $t_{u+1}^r = V_i$. We have $\rho_t^r(1) = \text{Corr}(V_1, V_i) = 0$.
5. The arrivals a_{u-1}^r and a_{u+1}^r both come from the same source i , $i \neq 1$. This is exactly the same as Case 4. Again $t_u^r = V_1$ and $t_{u+1}^r = V_i$. $\rho_t^r(1) = 0$.

Hence under the first two cases $\rho_t^r(1) = \rho_t(1)$, while under the last three cases $\rho_t^r(1) = 0$. Let the minimum forward recurrence time from all other sources be

$V_{min} = \min_{\{i=2\dots r\}} V_i$. Case 1 and 2 occur when a_u^r and a_{u+1}^r both come from source 1. This implies $t_{1,j+1} < V_{min}$. Therefore,

$$\Pr(\text{Case 1 or 2}) = \Pr(t_{1,j+1} < V_{min}).$$

Thus, we have

$$(16) \quad \rho_t^r(1) = \Pr(t_{1,j+1} < V_{min}) \text{Corr}(t_{1,j}, t_{1,j+1} | t_{1,j+1} < V_{min}).$$

It remains to solve for both terms on the right hand side of Equation 16. Based on the density of a forward recurrence time from one source process (Cox, 1962), we obtain the density for V_{min} , the minimum of $r - 1$ forward recurrence times. Since $t_{1,j+1}$ is Weibull with parameters λ and α , and independent of V_{min} , we have

$$P(t_{1,j+1} < V_{min}) = 1 - (r - 1)\alpha \int_0^\infty e^{-2y^\lambda/\psi} \left(1 - \alpha \int_0^y e^{-x^\lambda/\psi} dx\right)^{r-2} dy,$$

where $\psi = \left(\alpha\Gamma(1 + \frac{1}{\lambda})\right)^{-\lambda}$.

Next we approximate the joint density of two consecutive interarrival times from the same source process, t_j and t_{j+1} , by converting the joint density of their Gaussian images z_j and z_{j+1} , $z_j = Z^{-1}(W(t_j))$, and ignoring the correlation between t_j and t_{j+1} in the Jacobian matrix. Let $\rho = \rho_z(1) = \text{Corr}(z_j, z_{j+1})$. We have the following approximate density, up to a normalizing factor:

$$\begin{aligned} f_W(t_j, t_{j+1}) &\propto \frac{1}{\sqrt{1 - \rho^2}} \exp\left\{\frac{-1}{2 - 2\rho^2} [\rho^2 (Z^{-1}(W(t_j)) + Z^{-1}(W(t_{j+1})))^2 \right. \\ &\quad \left. + 2(\rho^2 - \rho)Z^{-1}(W(t_j))Z^{-1}(T(t_{j+1}))]\right\} \\ &\quad \times \frac{\lambda^2}{\psi^2} (t_j t_{j+1})^{\lambda-1} \exp\{-(t_j^\lambda + t_{j+1}^\lambda)/\psi\}. \end{aligned}$$

Using $f_W(t_j, t_{j+1})$ combined with $P(t_{1,j+1} < V_{min})$, $\text{Corr}(t_{1,j}, t_{1,j+1} | t_{1,j+1} < V_{min})$ can be calculated.

There are two ways to compute $\rho_t^r(1)$, either using Equation 16 or directly using the approximate joint density $f_W(t_j, t_{j+1})$, because $f_W(t_j, t_{j+1})$ can be applied to two consecutive interarrival times t_u^r and t_{u+1}^r in the superposition process as well. This provides a numerical method for us to find $\rho_z^r(1)$ from a given $\rho_t^r(1)$.

We first obtain the estimate $\hat{\rho}_t^r(1)$ using Equation 16. With $\hat{\lambda}(r)$ obtained using Equation 14 and $\hat{\alpha}(r) = r\alpha$, we evaluate $f_W(t_j, t_{j+1})$ over a grid of potential $\rho_z^r(1)$ values. For each $\rho_z^r(1)$ value we compute the corresponding $\rho_t^r(1)$ directly using $f_W(t_j, t_{j+1})$. The estimate $\hat{\rho}_z^r(1)$ is the one that provides the closest match to $\hat{\rho}_t^r(1)$ obtained from Equation 16. Then we apply Equation 15 to have an estimate of $\hat{\theta}(r)$.

References.

- Patrice Abry, Richard Baraniuk, Patrick Flandrin, Rudolf Riedi, and Darryl Veitch. The multiscale nature of network traffic: Discovery, analysis, and modelling. *IEEE Signal Processing Magazine*, 19:28–46, 2002.
- M. Ashoura and T. Le-Ngoc. Priority queuing of long-range dependent traffic. *Computer Communications*, 31:3954–3963, 2008.
- A. Baiocchi and F. Vacirca. Tcp fluid modeling with a variable capacity bottleneck link. In *Proceedings of IEEE Infocom*, pages 1046–1054, Anchorage, AK, 2007.
- N. Bäuerle and U. Rieder. Optimal control of single-server fluid networks. *Queueing Systems: Theory and Applications*, 35:185–200, 2000.
- P. Belottia, A. Caponeb, G. Carellob, and F. Malucelli. Multi-layer mpls network design: The impact of statistical multiplexing. *Computer Networks*, 52:1291–1307, 2008.
- J. Beran. *Statistics for Long-Memory Processes*. Chapman & Hall, 1994.
- J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. Internet traffic tends toward poisson and independent as the load increases. In C. Holmes, D. Denison, M. Hansen, B. Yu, and B. Mallick, editors, *Nonlinear Estimation and Classification*, pages 83–109. Springer, New York, 2002.
- G. Carofiglia, M. Garetto, E. Leonardi, A. Tarello, and M.A. Marsan. Beyond fluid models: Modelling tcp mice in ip networks under non-stationary random traffic. *Computer Networks*, 51:114–133, 2007.
- W.S. Cleveland, C.L. Mallows, and J.E. McRae. Ats methods: Nonparametric regression for non-gaussian data. *Journal of the American Statistical Association*, 88:821–835, 1993.
- D. R. Cox. *Renewal Theory*. Methuen and Co. Ltd., London, 1962.
- Istvan Csabai. $1/f$ Noise in Computer Network Traffic. *Journal of Physics A-Mathematical and General*, 27:L417–L421, 1994.
- T.D. Dang, S. Molnar, and I. Maricza. Some results on multiscale queueing analysis. In *Proceedings of the 10th International Conference on Telecommunications ICT*, pages 1631–1638, Papeete, French Polynesia, 2003.
- F. M. de Pereira, N. L. S. da Fonseca, and D. S. Arantes. On the performance of generalized processor sharing servers under long-range dependent traffic. *Computer Networks*, 40:413–431, 2002.
- R. L. Dobrushin and P. Major. Non-central limit theorems for non-linear functionals of gaussian fields. *Probability Theory and Related Fields*, 50:27–52, 1979.
- P. Doukhan, G. Oppenheim, and M. Taqqu. *Theory and Applications of Long-Range Dependence*. Birkhauser Boston, 2003.
- N. G. Duffield. Economies of scale in queues with sources having power-law large deviations scalings. *Journal of Applied Probability*, 33:840–857, 1996.
- Ashok Erramilli, Onuttom Narayan, and Walter Willinger. Experimental Queueing Analysis with Long-Range Dependent Packet Traffic. *IEEE/ACM Transactions on Networking*, 4:209–223, 1996.
- Ashok Erramilli, Matthew Roughan, Darryl Veitch, and Walter Willinger. Self-similar traffic and network dynamics. In *Proceedings of the IEEE*, pages 800–819, 2002.
- A. Feldmann, A. A. Gilbert, and W. Willinger. Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic. In *ACM SIGCOMM*, pages 42–55, 1998a.
- A. Feldmann, A.C. Gilbert, W. Willinger, and T.G. Kurtz. The changing nature of network traffic: Scaling phenomena. *SIGCOMM Comput. Commun. Rev.*, 28:5–29, 1998b.
- Daniel R. Figueiredo, Benyuan Liu, Vishal Misra, and Don Towsley. On the Autocorrelation Structure of TCP Traffic. *Computer Networks*, 40:339–361, 2002.
- Chuck Fraleigh, Fouad Tobagi, and Christophe Diot. Provisioning IP Backbone Networks to Support Latency Sensitive Traffic. In *IEEE INFOCOM*, 2003.
- J.B. Gao and I. Rubin. Multiplicative Multifractal Modeling of Long-Range-Dependent Network

- Traffic. *International Journal of Communications Systems*, 14:783–201, 2001a.
- J.B. Gao and I. Rubin. Superposition of Multiplicative Multifractal Traffic Streams. In *Proceedings of ICC2000*, 2001b.
- J. Geweke and S. Porter-Hudak. The estimation and application of long memory time series models. *Journal of Time Series Analysis*, 4:221–238, 1983.
- A.C. Gilbert, W. Willinger, and A. Feldmann. Scaling analysis of random cascades, with applications to network traffic. *IEEE Trans. on Information Theory*, 45(3):971–991, 1999.
- Wei-Bo Gong, Yong Liu, Vishal Misra, and Don Towsley. Self-Similarity and Long Range Dependence on the Internet: A Second Look at the Evidence, Origins and Implications. *Computer Networks*, 48:377–399, 2005.
- Y. Gu, Y. Liu, and D. Towsley. On integrating fluid models with packet simulation. In *Proceedings of IEEE Infocom*, pages 2856–2866, Hong Kong, 2004.
- J. Hannig, J. S. Marron, and R. H. Riedi. Zooming Statistics: Inference Across Scales. *Journal of the Korean Statistical Society*, 30:327–345, 2001.
- D. P. Heyman and T. V. Lakshman. What Are the Implications of Long-Range Dependence for VBR-Video Traffic Engineering? *IEEE/ACM Transactions on Networking*, 4:301–317, 1996.
- J. R. M. Hosking. Fractional Differencing. *Biometrika*, 68(1):165–176, 1981.
- <http://caida.org/tools>. The CAIDA Tools Site.
- <http://www.endace.com/>. Endace DAG.
- <http://www.wand.net.nz/wits/catalogue.php>. WITS: Waikato Internet Traffic Storage.
- C.M. Hurvich, E. Moulines, and P. Soulier. The fex estimator for potentially non-stationary linear time series. *Stochastic Processes and Their Applications*, 97:307–340, 2002.
- Hao Jiang and Constantinos Dovrolis. Why is the Internet traffic bursty in short time scales. In *Sigmetrics*, pages 241–252. ACM Press, 2005.
- Thomas Karagiannis, Mart Molle, Michalis Faloutsos, and Andre Broido. A nonstationary poisson view of internet traffic. In *Proceedings of IEEE INFOCOM*, 2004.
- G. Kesidis. *An Introduction to Communication Network Analysis*. Wiley-IEEE Press, 2007.
- C. Kiddle, R. Simmonds, C. Williamson, and B. Unger. Hybrid packet/fluid flow network simulation. In *Proceedings of the Seventeenth Workshop on Parallel and Distributed Simulation*, pages 143–152, Washington, DC, USA, 2003.
- H.S. Kim and N.B. Shroff. Loss probability calculations and asymptotic analysis for finite buffer multiplexers. *IEEE/ACM Transactions on Networking*, 9:755–768, 2001.
- Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.
- N.X. Liu and J.S. Baras. Statistical modeling and performance analysis of multi-scale traffic. In *Proceedings of IEEE Infocom*, 2003.
- Y. Liu, F.L. Presti, V. Misra, D.F. Towsley, and Y. Gu. Scalable fluid models and simulations for large-scale ip networks. *ACM Trans. Model. Comput. Simul.*, 14:305–324, 2004.
- Masao Masugi and Takehisa Takuma. Multi-fractal analysis of IP-network traffic for assessing time variations in scaling properties. *Physica D: Nonlinear Phenomena*, 225:119–126, 2007.
- K. Maulik and S. Resnick. Small and large time scale analysis of a network traffic model. *Queueing Systems*, 43:221–250, 2003.
- T. Mikosch, S. Resnick, H. Rootzen, and A. Stegeman. Is network traffic approximated by stable levy motion or fractional brownian motion? *Annals of Applied Probability*, 12(1):23–68, 2002.
- Ilkka Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.
- K. Park, G. Kim, and M. Crovella. On the Effect of Traffic Self-similarity on Network Performance. In *Proceedings SPIE Intl. Conf. Perf. and Control of Network Systems*, 1997.
- Vern Paxson. Fast Approximate Synthesis of Fractional Gaussian Noise for Generating Self-Similar Network Traffic. *Computer Communications Review*, pages 5–18, 1997.
- Vern Paxson and Sally Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM*

- Transactions on Networking*, 3:226–244, 1995.
- L. L. Peterson and B. S. Davie. *Computer Networks: A Systems Approach*. Morgan Kaufmann, 1999.
- S. Resnick, A.C. Gilbert, and W. Willinger. Wavelet analysis of conservative cascades. *Bernoulli*, 9(1):97–135, 2003.
- Vinay J. Ribeiro, Zhi-Li Zhang, Sue Moon, and Christophe Diot. Small-time scaling behavior of Internet backbone traffic. *Computer Networks*, 48(3):315–334, 2005.
- V.J. Ribiero, R. H. Riedi, and R. G. Baraniuk. Multiscale queuing analysis. *IEEE/ACM Transactions on Networking*, 14(5):1005–1018, 2006.
- R. Riedi and J.L. Vehel. Multifractal properties of TCP traffic: a numerical study. Technical Report 3129, INRIA Rocquencourt, France, 1997.
- Rudolf H. Riedi. Multifractal processes. In Paul Doukhan, George Oppenheim, and Murad S. Taqqu, editors, *Theory and Application of Long-Range Dependence*, pages 625–715. Birkhauser, Basel, 2002.
- Rudolf H. Riedi, Matthew S. Crouse, Vinay J. Ribeiro, and Richard G. Baraniuk. A Multifractal Wavelet Model with Application to Network Traffic. *IEEE Transactions on Information Theory*, 45(3):992–1019, 1999.
- D. A. Rolls, G. Michailidis, and F. Hernández-Campos. Queueing analysis of network traffic: Methodology and visualization tools. *Computer Networks*, 48(3):447–473, 2005.
- Matthew Roughan and Darryl Veitch. Some remarks on unexpected scaling exponents. *SIGCOMM Computer Communications Review*, 37:71–74, 2007.
- W.R. Stevens. *TCP/IP Illustrated, Volume 1: The Protocols*. Addison-Wesley, 1994.
- Stilian Stoev, Murad S. Taqqu, Cheolwoo Park, and J.S. Marron. On the Wavelet Spectrum Diagnostic for Hurst Parameter Estimation in the Analysis of Internet Traffic. *Computer Networks*, 48:423–445, 2005.
- D. Surgailis. Long-range dependence and appell rank. *The Annals of Probability*, 28:478–497, 2000.
- M. Taqqu. Weak convergence to fractional brownian motion and to the rosenblatt process. *Probability Theory and Related Fields*, 31:287–302, 1975.
- M. Taqqu. Law of the iterated logarithm for sums of non-linear functions of gaussian variables that exhibit a long range dependence. *Probability Theory and Related Fields*, 40:203–238, 1977.
- M.S. Taqqu, V. Teverovsky, and W. Willinger. Is network traffic self-similar or multifractal. *Fractals*, 5:63–73, 1997a.
- M.S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, 27:5–23, 1997b.
- Darryl Veitch, Nicolas Hohn, and Patrice Abry. Multifractality in TCP/IP Traffic: the Case Against. *Computer Networks*, 48:293–313, 2005.
- A. Veres and M. Boda. The Chaotic Nature of TCP Congestion Control. In *IEEE INFOCOMM*, pages 1715–1723, 2000.
- W. Willinger, M.S. Taqqu, and A. Erramilli. A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks. In F. P. Kelly, S. Zachary, and I. Ziedins, editors, *Stochastic Networks: Theory and Applications*, pages 339–366. Clarendon Press (Oxford University Press), Oxford, 1996.
- W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling Phenomena in the Internet: Critically Examining Criticality. *Proceedings of The National Academy of Sciences of The United States of America*, 99(3):2573–2580, 2002.
- Walter Willinger, Murad S. Taqqu, Will E. Leland, and Daniel V. Wilson. Self-Similarity in High-Speed Packet Traffic: Analysis and Modeling of Ethernet Traffic Measurements. *Statistical Science*, 10:67–85, 1995.
- Jin Yuan, Yong Ren, and Xiuming Shan. Self-Organized Criticality in a Computer Network Model. *Physical Review E*, 61(2):1067–1071, 2000.

DEPARTMENT OF MATHEMATICS,
XAVIER UNIVERSITY OF LOUISIANA,
NEW ORLEANS, LA 70125, USA
E-MAIL: danders2@xula.edu

DEPARTMENT OF STATISTICS,
PURDUE UNIVERSITY,
WEST LAFAYETTE, IN 47907, USA
E-MAIL: wsc@stat.purdue.edu

DEPARTMENT OF STATISTICS,
PURDUE UNIVERSITY,
WEST LAFAYETTE, IN 47907, USA
E-MAIL: xbw@stat.purdue.edu